

EPISTEMIC INTERNALISM: AN EXPLANATION AND DEFENSE

Brent James Charles Madison

Department of Philosophy

University College London

A thesis submitted in fulfilment
of the requirements for the degree of
Ph.D.

University of London
March 2008

UMI Number: U591622

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U591622

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

ABSTRACT OF THESIS

DECLARATION:

I, Brent James Charles Madison, confirm that the work presented in this thesis is my own.

Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

DATE:

Mar. 7, 08

SIGNATURE:

UCL REGISTRY

Department of Student Administration

Student Records

UCL, Gower Street London WC1E 6BT

Tel: +44 (0)20 7679 7005/6 Fax: +44 (0)20 7679 7920

researchdegrees@ucl.ac.uk



ABSTRACT OF THESIS

Full name as registered on Portico. Please type surname in BLOCK CAPITALS

Full Name: **Brent James Charles MADISON**

Student Number: **49024780**

Degree for which you are entering (please tick): PhD ☒ MPhil ☐ MD(Res) ☐ MD ☐ MS ☐

Title of thesis as agreed by your supervisor:

Epistemic Internalism: An Explanation and Defense

Please type your abstract below. It must be limited to one side of single-spaced type. The Library will circulate the sheet to the publisher of the Index to Theses.

What does it take for a positive epistemic status to obtain? I argue throughout my thesis that if a positive epistemic status obtains, this is not a brute fact. Instead, if for example a belief is justified, it is justified in virtue of some further condition(s) obtaining. A fundamental topic in epistemology is the question of what sorts of factors can be relevant to determining the positive epistemic status of belief. *Epistemic Internalism* holds that these factors must be “internal” (in a sense that needs to be specified). *Epistemic Externalism* is the denial of internalism. My thesis is an explanation and defense of an internalist theory of epistemic justification.

The central claim of my thesis is that something is “internal” in this sense only if it is, or can easily be, the object of the agent’s conscious awareness. By considering key cases, I show that without an awareness requirement on justification, the subject cannot avoid what I call the Subject’s Perspective Objection. In developing this objection I examine and respond to an argument against the awareness requirement which claims that such awareness either leads to a vicious regress of requiring higher order beliefs of increasing complexity (if *any* beliefs are to be justified), or else requiring such awareness is unmotivated. This regress is generated because it is assumed that the relevant kind of awareness must be *doxastic*. My solution invokes what I call ‘strong non-doxastic awareness’ that grounds non-inferential justification and thereby avoids this dilemma, while meeting the Subject’s Perspective Objection.

I also argue that external factors, such as the reliability (actual or conditional) of the mechanism supporting the belief, are not necessary for justification. I argue for this conclusion by comparing what constitutes justified belief in the actual world with one’s counterpart in a ‘demon world’. I argue that this intuition, correctly interpreted, counts in favour of internalism. As I explain, many philosophers moved by arguments presented by externalists about mental content deny that such a case is possible. In opposition I argue that the awareness requirement remains substantially unaffected, no matter which view of content turns out to be correct. What is key is that the two worlds are completely subjectively indistinguishable from each other for those who inhabit them in all the ways of which they are consciously aware.

If neither the obtaining of truth nor reliability is necessary for epistemic justification, what makes justification genuinely epistemic? In the final chapter I argue against recent work that assimilates justification with knowledge, as well as for a positive account of the truth connection. As to the former question, I defend the orthodoxy that they are distinct epistemic statuses; as to the latter, I argue that the connection between justification and truth is *conceptual*. That is, epistemic justification is epistemic because it turns on evidence; evidence is epistemic because it is conceptually linked with truth. Epistemic justification, therefore, is conceptually linked with the truth (via evidence), which is what makes it distinctively epistemic.

In short, this thesis is an explanation and defense of an internalist theory of epistemic justification. The central claim of the thesis is that something is “internal” in this sense only if it is, or can easily be, the object of the agent’s conscious awareness. I argue that conscious awareness is a necessary condition of epistemic justification obtaining, and that factors external to consciousness play no justificatory role.

This thesis is dedicated to the women in my life:

My Mother Vickie,

My Wife Noni

and

My Daughter Regan

ACKNOWLEDGEMENTS

First and foremost I would like to thank those who have taught me philosophy. It was Michael Vertin and Doug Hutchinson who first introduced me to philosophy as an undergraduate at the University of Toronto, and who inspired me to pursue it.

At King's College London I would particularly like to thank Anthony Savile, Chris Hughes and David Galloway for tutoring me in preparation for my M.Phil exams that served as my introduction to analytic philosophy. Also at King's, special thanks to Peter Byrne who supervised my M.Phil thesis and got me thinking about issues in epistemology.

At UCL I was also fortunate enough to have had excellent tutors who helped me discover, develop and deepen my philosophical interests. During the year he supervised me, it was Quassim Cassam who helped me identify that it was an interest in defending epistemic internalism that lay behind my inchoate intuitions about 'epistemic rationality'. Over the next year Mike Martin helped me develop and refine my views through intense engagement with my work. More recently, Paul Snowdon has been excellent in helping me finish the thesis and has consistently pressed me to iron out the details of my position.

In London I have benefited from numerous discussions with my peers, but I would especially like to thank Rod Farningham, Johan Gersel, Adil Mawani, Alex Steinberg and Lee Walters. Since arriving in Sydney, Stephen Hetherington has been particularly generous in providing feedback on my work.

My mother Vickie deserves special thanks for just about everything, but in particular, for being a formidable copy-editor and proofreader who has valiantly suffered through virtually every page I have ever written since I was in junior high school. This thesis is dedicated to her, as well as to my wife Noni and our daughter Regan. All three, in their own ways, have been enormous sources of joy, support and encouragement.

TABLE OF CONTENTS

| | |
|--|----|
| Declaration | 2 |
| Abstract | 3 |
| Dedication | 4 |
| Acknowledgements | 5 |
| Introduction | 8 |
| Chapter 1. Preliminaries: On The Nature of Epistemic Justification | 18 |
| 1.1. What Kinds of Epistemic Evaluation Does the Internalism/Externalism Distinction Apply To? | 18 |
| 1.2. Basic Features of Justification | 22 |
| 1.3. Doxastic v. Propositional Justification and The Basing Relation | 26 |
| 1.4. Some Ways of Construing the Internalism/Externalism Distinction | 46 |
| 1.5. The Structure of Justification and the Existence of Immediate Justification | 52 |
| Chapter 2. The Awareness Requirement and Its Nature | 68 |
| 2.1. On The Necessity of Awareness | 68 |
| 2.2. Mentalism | 90 |
| 2.3. On the Nature of Awareness: Responding To A Dilemma for Internalism | 97 |

| | |
|--|------------|
| Chapter 3. The Case of the New Evil Demon: Is (Actual) Reliability Necessary for Justification? | 119 |
| 3.1. Introduction | 119 |
| 3.2. What is the “New Evil Demon” Intuition? | 122 |
| 3.3. Content Externalism | 125 |
| 3.4. A Reply to Williamson on the Incompatibility of Epistemic Internalism and Content Externalism | 134 |
| 3.5. Other Reasons to Reject the New Evil Demon Intuition: McDowell’s Epistemological Disjunctivism | 139 |
| Chapter 4. The Truth-Connection: What It Is Not – Reliability Redux | 155 |
| 4.1. Introduction | 155 |
| 4.2. The Subject’s Role and the Truth Connection: Normal World Reliabilism | 159 |
| 4.3. Peacocke On Entitlement and Truth-Conduciveness | 165 |
| 4.4. Peacocke’s <i>A Priori</i> Arguments Against Scepticism | 173 |
| Chapter 5. A Positive Account of the Truth Connection, or, What Makes Justification Epistemic (and Is Justification Distinct From Knowledge)? | 188 |
| 5.1. Introduction | 188 |
| 5.2. On Epistemic Goals / Ends | 190 |
| 5.3. Is Justification Nothing More, and Nothing Less, Than Knowledge? | 198 |
| 5.4. The Conceptual Connections Between Justified Belief and Truth | 224 |
| Conclusion | 236 |
| Bibliography | 241 |

INTRODUCTION:

This thesis concerns the nature of epistemic justification. But despite having tried to offer a well-explained, cogent position on certain aspects of the nature of justification, one might well object to many of the remarks in this thesis as utterly question-begging. But such a complaint changes the topic from the one that I am addressing here. Obviously, the approach taken throughout this thesis will not convince those who are sceptical about whether our beliefs are (for the most part) justified, or think that in order to be justified we need to have a non-circular (and therefore, presumably *a priori*) demonstration that proves that we are justified in believing much of what we do. But again, that is to engage in a different debate, one outside the scope of this project, since I will be using what can be called ‘Particularism’ as a methodological constraint in exploring the nature of epistemic justification, as I will explain below.

A fundamental problem in epistemology is generated by the conjunction of two questions: what do we know, and how do we decide, in a particular case, whether we know? As Chisholm puts the latter question, “what are the *criteria* of knowing?” (Chisholm, 1966, p. 56) Of course, perhaps there will prove to be no such criteria. But a working hypothesis to be tested throughout this thesis is that we can formulate, at least partially, illuminating conditions for the instantiation of certain epistemic properties.

Chisholm’s idea is that if we could tell what we know, we could go about trying to formulate the criteria of knowing. Similarly, if we had such criteria, we could determine what, if anything, we know. A sceptical position is to deny that we have an answer to either of these questions, so that we cannot know what, if anything, we know,

and we have no way of deciding in a particular case if we genuinely know it. But there are alternatives for non-sceptics.

Those who claim to have an intuitive idea of what we know, are called ‘Particularists’, and they seek to answer the criteria question by giving an account of what all the intuitive particular cases of knowledge have in common that makes them instances of *knowledge*, rather than something else. On the other hand, those who claim to have general criteria of knowledge are sometimes called ‘Methodists’. They try to determine what (if anything), and to what extent, matches the criteria of knowledge that they propose. This sort of problem is not unique to knowledge. I will be presupposing Particularism throughout as a methodological constraint on approaching the theory of epistemic *justification*. That is, given that I share an intuitive idea of what we are justified in believing, I will seek to give an account of what all the intuitive particular cases of justification have in common that make them instances of *justification*, rather than something else.

I hold that we are justified in believing all sorts of things, e.g. I am justified in believing that there is a desk before me, that the room is now less than 50 degrees C., etc. This project *presupposes* that we are epistemically justified in believing most of what we take ourselves to be justified in believing. Given that we obviously *are* justified in believing many things, can we elucidate what it is to be justified? Part of my answer to this question is expressed by arguing that the Awareness Requirement introduced in Chapter 2 is a necessary condition of justification obtaining. Given this approach, it is no objection to my account to merely point out that it would not convince the sceptic or that it presupposes that we have justified beliefs. These things are true, but they are beside

the point. Again, the question I will be addressing is that, given that we have justified beliefs, what can we say about the nature of justification?

I will argue, among other things, for the indispensability of conscious awareness to epistemic justification. In effect I will be defending a version of what has been labelled, not particularly helpfully, ‘internalism’ in epistemology.¹

The so-called internalism/externalism debate should be of interest in epistemology since it addresses one of the most fundamental questions in the discipline: what is the basic nature of epistemic justification? The question of what *kinds* of factors can be relevant to determining the positive epistemic status of belief seems to be about as central as a meta-epistemological question could be.² One might think that it is only with firm grasp of meta-epistemology that one can properly approach questions of *applied* epistemology, such as what, if anything, are we justified in believing, and how is such justification possible? It is only once we know what we are looking for that we should start our search for it.

Obviously we are justified in believing lots of different *kinds* of propositions, and how we acquire these beliefs and their justification varies as well. Certain kinds of

¹ While there are obviously epistemologies in the history of philosophy that can be classified as either ‘internalist’ or ‘externalist’ in nature, the labels themselves were not used until relatively recently. David Armstrong (1973) introduced what he called ‘externalist’ theories of knowledge as those that hold that knowledge is a natural or law-like connection between a subject’s true belief and the state of affairs that make it true. This is what has become known as his ‘Thermometer Model of Knowledge’.

However, since it is not clear what role, if any, epistemic *justification* plays in Armstrong’s epistemology, it is not clear to what extent his concept of the epistemically external meshes with the current debate between internalists and externalists who speak explicitly in terms of justification. Therefore it was perhaps not until the exchange between Laurence Bonjour (1980) and Alvin Goldman (1980) who offered ‘internalist’ and ‘externalist’ theories of epistemic justification that these labels were first used in their current sense.

² For the distinction between meta- and applied epistemology, and the case for beginning our epistemological enquiry with the meta-epistemological, see Fumerton (1995) Chapter 1.

beliefs, such as memory beliefs, or *a priori* beliefs, might be thought to pose special problems for the kind of internalism I go on to defend here.

For example, internalists tend to hold that so long as a justifier is available to the subject either immediately or upon introspection, it can serve to justify beliefs. Here Alvin Goldman has objected with what he calls “the Problem of Forgotten Evidence”. (See Goldman 2001) The fact is that many beliefs are ones which were based on adequate grounds available to the subject at the time the belief was formed, but as time lapses, the grounds are often forgotten. As Goldman points out, how properly the belief was originally formed is irrelevant to many internalists. So, at this later time when the grounds are forgotten, the worry is that the would-be grounded belief is unjustified (which leads to scepticism) because the justifying grounds are no longer available to the subject, even upon introspection.

While the Problem of Forgotten Evidence and the justification of memory beliefs, as well as *a priori* beliefs, beliefs based on testimony, etc. are interesting cases that a complete theory of epistemic justification should ultimately address, I will not be able to address these problems in this thesis. Instead they must be left for further research. I will for the most part limit my scope to beliefs in contingent, empirical propositions about the external world based on sense perception.

The justification of perceptual belief is a central case of justification. Furthermore, it is *acquiring* justified belief that is the central case of epistemic justification, not its mere retention. I place special emphasis on the acquisition of belief, rather than its retention, because it is conceptually more fundamental: one cannot retain a belief unless one has acquired it; but one can acquire a belief but fail to retain it (perhaps

because one has poor memory). In the future I intend to extend the account I develop here of the justification of perceptual belief to other potentially problematic areas. In each case I intend to offer a response that explains how the type of justification is possible that is neither externalist, nor concessive to the sceptic, in a way that is consistent with epistemic internalism and the awareness condition.³

But first, a concern. By requiring the capacity for conscious awareness as a necessary condition for being epistemically justified, will my account of justification be too intellectually demanding, thereby barring most if not all people from holding justified beliefs? People vary in intellects and abilities, and accordingly, people's degree of justification for the very same belief will vary to reflect that. If a person's ability allows him to recognize an *a priori* or evidential insight as my account requires, he can use it to justify a corresponding belief. Not being able to achieve maximal justification (as presumably no fallible epistemic agents, such as humans, can hope to achieve) does not mean that an agent cannot have a justified belief to any degree whatsoever. Since justification is almost universally thought to come in degrees, as I will explain in Chapter 1, it is no objection to a theory if it allows that people of varying levels of intellect are capable of achieving different levels of justification.

I think there is some credence to charges of over-intellectualization if it is directed at a theory of knowledge generally. If young children and some animals do indeed have knowledge (which I take to be an open question – and not one to which the answer is obvious), over-intellectualization is something to be avoided. However, while over-intellectualization is by definition a bad thing, intellectualization *simpliciter* is not. It is

³ For examples of internalist responses to the Problem of Forgotten Evidence, see Bonjour (2001) and Feldman and Conce (2001). For an example of an internalist theory of *a priori* justification, see Bonjour (1998).

true that understanding justification normatively, in terms of conscious awareness, may require a capacity for critical self-reflection. Fortunately, these are just the kind of capacities normal adult humans have. We are intellectual beings, or as Aristotle would have it, rational animals. If young children or some animals lack the capacity for conscious awareness or critical reflection on what they believe, they cannot have justified beliefs.

But that is not to say that their beliefs are unjustified. To be unjustified is to have the capacity for rationality but to fall short of meeting epistemic standards. If one lacks a capacity and accompanying standards, one cannot be faulted for not meeting them. For example, I lack the capacity for unaided flight. But it would be absurd to fault me for being unable to fly. It would be equally absurd to think that evaluations of unaided flight even apply to me. Rather, children, animals and those who lack the capacity for critical self-reflection are *a-justificational* – that is, the concept of justification does not apply to them, making evaluations of justified belief inapplicable. What I am considering here is the nature of epistemic justification, and those individuals and their mental states of which it is a property.

An Overview:

In Chapter 1 I address some preliminary issues and defend some fundamental distinctions concerning the nature of epistemic justification. For instance, I outline some of the basic features of justification, distinguish doxastic from propositional justification and consider how best to understand the epistemic basing relation. Specifically, what role (if any) does causation play in properly basing a belief on a reason? I also consider the structure of justification and argue for the existence of immediate justification, that is, the justification that does not come from one's justification to believe other propositions. In effect I defend Foundationalism against its rivals (e.g., Coherentism). In addition I also explore ways of construing the internalism/ externalism distinction, and consider to what kinds of epistemic evaluation the distinction applies. I conclude the chapter by outlining the egocentric or first person perspective and note the kinds of constraints that such a perspective puts on an adequate theory of epistemic justification.

In Chapter 2 I motivate the need for *awareness* as a necessary condition of epistemic justification obtaining. I capture the kind of awareness justification requires with the so-called Awareness Requirement, which is the hallmark of the kind of internalism I am defending. The condition can be stated as follows:

Awareness Requirement:

S's is justified in believing that *p* only if

- i) There is something, X, that contributes to the justification of S's believing that *p*; and
- ii) For all X that contributes, S is aware (or potentially aware) of X.

I argue for this requirement by drawing attention to our considered judgements about cases. I consider two kinds of cases: first, I clarify and draw out the implications of cases of clairvoyance and other cases of unusual but reliable cognitive faculties. Secondly, I show that even many of those who claim to reject such awareness requirements implicitly appeal to them to motivate their own accounts of justification. By considering the arguments for “Mentalism”, I show that, unless an awareness requirement is presupposed, the cases adduced by Mentalists are of no intuitive force. In turn, by considering just these cases I aim to establish that the awareness requirement is necessary for justification.

In the latter half of chapter 2 I provide an account of the nature of this awareness. In doing so I examine and respond to an argument against the awareness requirement that we might call Bergmann’s Dilemma. My solution invokes what I call ‘non-doxastic strong awareness’ in a way that avoids the regresses that Bergmann advances. The overall aim of Chapter 2 is to establish that factors external to the subject’s consciousness (such as the reliability of the process that gives rise to the belief in question) are not *sufficient* for justification. Awareness is required.

In Chapter 3 I consider whether external factors are nevertheless *necessary* for justification to obtain. I argue that they are not. This is done by considering the case of the New Evil Demon. After explaining what this case is, I explain how it supports Fallibilism, the Awareness Requirement, as well as the claim that the *actual* reliability of the belief-forming mechanism used in a world is not necessary for epistemic justification to obtain. I then defend this position against recent influential arguments advanced by Timothy Williamson and John McDowell, respectively.

Some connection to truth is what makes a given thing epistemic. In Chapter 4 I investigate proposals that conceive of the relation between justified belief and truth as one of truth-conduciveness. I argue that, given internalism about justification, either in the sense I defend in Chapter 2 involving the Awareness Requirement, or in a weaker sense that Christopher Peacocke endorses, the truth-relation cannot be characterized in terms of truth-conduciveness or reliability. This is not just to reiterate a conclusion of Chapter 3 that justification need not be *actually* truth-conducive. I argue further that the truth-connection cannot even be one of truth-conduciveness in a *conditional* mode. Having dispatched truth-conducive accounts of the truth-connection, I am in a position to develop a positive account of the truth-connection in Chapter 5.

Epistemic justification is different from other kinds of justification. Moral, aesthetic and pragmatic justification (assuming there are such things), seem to relate essentially to goodness, beauty and usefulness, respectively. What individuates *epistemic* justification from the other types is *some* connection to truth. So far we have seen what this connection is not: in Chapter 3 I argued that justification need not be *actually* truth-conducive. In Chapter 4 I argue further that the truth-connection cannot even be one of truth-conduciveness in a *conditional* mode.

In Chapter 5 I argue against recent work that assimilates justification to knowledge, as well as for a positive account of the truth connection. As to the former question, I defend the orthodoxy that they are distinct epistemic statuses; as to the latter, I argue that the connection between justification and truth is *conceptual*. That is, epistemic justification is epistemic because it turns on evidence; evidence is epistemic because it is conceptually linked with truth. Epistemic justification, therefore, is conceptually linked

with the truth (via evidence), which is what makes it distinctively epistemic. This account of the truth-connection is substantial enough to ensure that the kind of justification we are concerned with is genuinely epistemic, while still being consistent with the kind of epistemic Internalism I am advocating.

In short, this thesis is an explanation and defense of an internalist theory of epistemic justification. The central claim of the thesis is that something is “internal” in this sense only if it is, or can easily be, the object of the agent’s conscious awareness. I argue that conscious awareness is a necessary condition of epistemic justification obtaining, and that factors external to consciousness play no justificatory role.

CHAPTER ONE

PRELIMINARIES: ON THE NATURE OF EPISTEMIC JUSTIFICATION

1.1. What Kinds of Epistemic Evaluation Does the Internalism / Externalism Distinction Apply To?

Epistemic ascriptions are a familiar part of our daily lives. We say that a subject *knows* that *p*, or is not *justified* in believing that *p*, or is *rational* in believing that *p*. What these cases all have in common is that we are *evaluating* a person's belief from what we might call the *epistemic* perspective. We are judging that a certain epistemic property is present. More will be said on what the epistemic perspective is, but roughly, I will argue that it is one that concerns *truth*, rather than beauty, moral goodness, usefulness, etc. To credit a belief with a *positive epistemic status* is to evaluate that belief as an epistemic good.

But what does it take for positive epistemic status to obtain? I shall argue in this thesis that if positive epistemic obtains, this is not a brute fact. Instead, if a belief is, for example, justified, it is justified in virtue of some further condition(s) obtaining. A fundamental topic in epistemology, the one that will be the subject of this thesis, is the question of what sorts of factors can be relevant to determining the positive epistemic status of belief. What has been called epistemic internalism holds, as the label suggests, that all the relevant factors must (in a sense that needs to be specified) be "internal". Epistemic externalism is the denial of internalism. A central claim of this thesis will be that "internal" in this sense is to be understood as those things that are, or easily can be,

available to the agent's conscious awareness. So the relevant sense of 'internal' here will mean internal to the subject's first person perspective, or point of view. I will argue that such awareness is a necessary condition for certain positive epistemic statuses obtaining, specifically, epistemic justification.

While developing and defending an internalist account of epistemic *justification* will be my primary focus in this work, in principle the internalism/externalism distinction can be applied to any epistemically normative status of belief. This point seems to be sometimes overlooked. Simply because, say, knowledge is external, since truth is a necessary condition and it is uncontroversially on all accounts an external property, it does not follow that the concepts that it breaks down into, if any, cannot be internalist. Similarly, there may be kinds of epistemic evaluation that are internalist but may not be necessary conditions for knowledge, e.g. "Foley-Rationality", or epistemic blamelessness/responsibility, or reasonableness.⁴ This is worth noting because it may turn out that certain kinds of internalism or externalism may hold for some epistemic concepts, but not for all, and so generalizations across epistemology should be made with great care. Despite the plurality of concepts of epistemic evaluation, however, the most widely discussed in the internalism/externalism debate concerns epistemic justification and its nature.

It is often assumed that justification is a necessary condition for knowledge, and before the Gettier cases of 1963, many thought that justification was sufficient to make a

⁴ "Foley-Rationality" is an account of epistemic rationality advanced by Richard Foley. See Foley (1987) and (1993) for details.

true belief knowledge.⁵ Following Foley, however, I am not going to assume from the outset that justification is a necessary condition for knowledge, let alone even related to it in a straightforward way. (Foley, 2005) Following Plantinga, I will call that concept (the one that essentially has to do with knowledge in that it marks the difference between true belief and knowledge) Warrant. (Plantinga, 1993) This of course assumes that there is such a notion.⁶ If there is not, this will obviously pose problems for those seeking a reductive analysis of knowledge in terms of necessary and sufficient conditions. My focus on *justification* as an epistemic evaluation in its own right, however, will remain unaffected regardless of what, if anything, amounts to Warrant.

Using this terminology, justification can be (and has been) offered as a theory of Warrant, but it need not be. I aim to elucidate aspects of the theory of justification, drawing connections to knowledge if and only if they arise, rather than focusing primarily on knowledge, and trying to show that it necessarily must be fundamentally understood in terms of justification. In other words, I shall remain agnostic respecting the relationship,

⁵ This seems to be the best explanation for the relatively little said against the Tripartite Analysis of Knowledge (JTB) before Gettier's counterexamples to the sufficiency of the analysis, compared to the onslaught of rival analyses that followed. For a survey of these attempts see Shope (1983).

There is also a different reading of that history that seems to have been in currency in Oxford. According to some who were in Oxford at the time, it is not that they thought the Tripartite analysis was correct, but that they did not think it worthwhile publishing an article showing it was wrong, and so ignored it in silence.

⁶ Some challenge this presupposition. Richard Swinburne, for example, expresses skepticism about the entire project of attempting to provide an analysis of Warrant. He argues in the spirit of Wittgenstein that there may not be any necessary and sufficient conditions for knowledge since it is too imprecise a concept. (Swinburne, 1995, p. 417) The suggestion is that there may be no one thing that when added to true belief turns it into knowledge. Instead, our concept of knowledge gets its meaning through correct and consistent application to instances of knowledge, though it is not specifiable in terms of exact criteria. Swinburne suggests that his claim has good inductive support since no acceptable list of necessary and sufficient conditions that analyses what Warrant is have been produced. (Swinburne, 1995, p. 419) Still, it obviously does not follow from the fact that there has not been a unanimously accepted analysis of Warrant thus far that one is not correct, or that a correct analysis cannot be produced.

Timothy Williamson also argues knowledge is not true belief plus some third thing. Instead, he also argues the thesis that our concept of knowledge is unanalysable in terms of necessary and sufficient conditions. In fact, according to Williamson, knowledge is *sui generis* mental state. For details see Williamson (2000).

if any, between justification and knowledge. Otherwise, there is a risk of giving a non-neutral account of justification that builds in internalist/externalist elements.

Before analyzing what it would mean to say that justification is “internal”, it is necessary to determine how epistemologists tend to characterize the basic features of justification.

1.2. Basic Features of Justification:

In a paper on justification, “Concepts of Epistemic Justification”, William Alston identifies four features common to justification: the first is that justification applies to beliefs. (Alston, 1989a, p. 83) It is true that beliefs have traditionally been the propositional attitude of choice as far as epistemic justification is concerned. Perhaps this has been because belief is held to be the only propositional attitude that is a necessary condition for knowledge. But once we leave questions of knowledge aside, and focus on the theory of epistemic justification, it seems that limiting the application of justification to beliefs is unduly restrictive. Instead, we should hold that justification applies to our epistemic *attitudes*, of which belief is just one. Other attitudes such as disbelief⁷ or suspending belief are also subject to justification. (Feldman, 2002, p. 368)

The second feature of justification is that it is an evaluative concept, as opposed to a merely “factual” or descriptive concept. (Alston, 1989a, p. 83) Alston’s use of “evaluative” seems to be what might also be described as a normative concept. As Alston puts it, to say that one is justified in believing that *p*, “is to accord S’s believing with a positive evaluative status.” (Ibid.) Being justified, therefore, is an epistemic good.

A third and more elusive aspect of justification that I alluded to earlier is, as Alston says, “to do with a specifically *epistemic* dimension of evaluation.” (Ibid.) The use of the qualifier “epistemic” is needed to individuate epistemic justification from other sorts, such as moral, aesthetic, or pragmatic justification. The way Alston defines what is uniquely epistemic is in terms of what he thinks of as the epistemic goal. According to

⁷ This of course assumes that besides believing that *p*, and believing that $\sim p$, there is the distinct attitude of *disbelieving* that *p*. I am not committing myself to the existence of such an attitude as Feldman does, but rather to the claim that if there is such a state, it is an *epistemic attitude*, and as such, epistemic justification is a potential property of it.

him, that is “the aim [of] maximizing truth and minimizing falsity in a large body of beliefs.” (Ibid.) Alston admits that this is only a rough formulation and needs refinement.

However, a problem with this kind of goal that requires one to maximise two variables at once is that one tends to end up with a certain kind of undecideability. For example, which set of beliefs best meets Alston’s goal: *q*: a set of 10 true beliefs and 1 false belief; or *m*: a set of 20 true beliefs and 2 false beliefs. *q* contains fewer false beliefs, but *m* contains more truths. Which set maximizes truth and minimizes falsehood, and thus meets the proposed epistemic goal? It does not seem that one can say. Therefore, a single goal to maximize would perhaps be better, e.g. for each belief that one holds, believe it only if it is true. At any rate, I agree with Alston that *some* connection to truth is what makes a given thing epistemic. I will address this issue of the truth-connection in chapters 4 and 5. For the moment, however, I suggest we leave open what exactly the truth connection is, to avoid any stipulations that will prejudice substantive issues in the internalism/externalism debate.

The final feature of justification Alston notes is that justification comes in degrees. (Alston, 1989a, p. 84) This means that one can have more or less justification for a given proposition. For example, suppose I observe that the sun rises one morning, thereby giving me some justification for thinking it true that the sun rises every morning. Depending on the conditions under which I observe the sunrise, *inter alia*, my belief that the sun rises in the morning will have some *prima facie* justification. Presumably, the more times I see the sun rise in the morning, the *stronger* my justification for my general belief that the sun rises every morning. This seems to show that justification comes in

degrees. Still, that being so, for simplicity I will follow Alston in treating justification as an all-or-nothing matter unless otherwise noted.

Similarly, some philosophers speak of *belief* as coming in degrees. Some also speak of the degree of *credence* to which a proposition is held.⁸ It is unclear if credence and belief really are the same cognitive attitude.⁹ But if they are, or if belief does indeed come in degrees, for simplicity, I will treat belief as an all-or-nothing matter. Nothing I go to argue throughout this thesis turns on whether or not belief is gradable.

I think that these four features are all that can be said neutrally to characterize justification. Others disagree, holding that justification is essentially connected to knowledge, or presupposes particular substantive connections between justification and truth.

For example, James Pryor in attempting to “neutrally” characterize epistemic justification, writes, “I say that you have justification to believe *p* iff you are in a position where it would be epistemically appropriate for you to believe *p*.” (Pryor, 2005, p. 181) So far so good -- this is just affirming Alston’s second feature of justification, namely, that justification is a normative status, an epistemic good. In elaborating on what Pryor means by “epistemically appropriate”, he continues, “a position where *p* is epistemically likely for you to be true.” (Ibid.) What does “epistemically likely for you to be true mean”? Pryor does not say. If he means an objective probability of being true, his basic account of justification already presupposes a form of externalism. If he means a subjective, epistemic conception of likelihood, he may already be incorporating internalist elements. Accordingly, in the aim of presupposing as little as possible, all

⁸ See for example Kaplan (1996).

⁹ See Smithies (manuscript) for a defense of the claim that beliefs are psychologically real, but are not reducible to states of confidence.

should agree justification is a matter of “epistemic appropriateness” having the four features Alston mentions, but nothing more can be said that does not bias the debate, including tying justification essentially to knowledge or to particular conceptions of likeliness or truth-conduciveness. So, with a rough characterization of what constitutes justification, we are getting closer to being able to ask what it would mean to be internalist about justification.

1.3. Doxastic v. Propositional Justification and The Basing Relation:

Doxastic v. Propositional Justification:

It is important to note that not all epistemologists cast their theory of justification in terms of justified *belief* (doxastic justification). Sometimes their concern is with, as Pryor puts it, “whether you have justification for believing certain *propositions*—regardless of whether you actually do believe those propositions.” (emphasis in original) (Pryor, 2001, p. 104) I shall call this kind of justification propositional justification. Richard Feldman, on the other hand, draws the contrast in terms of “being justified in believing the proposition (since she has good reasons) but that she does not believe it justifiably (since she bases her belief on something other than good reasons).¹⁰ (Feldman, 2005, p. 274). Expressing Feldman’s distinction in our terms, he is saying that when one has good reasons one’s belief enjoys propositional justification, but one’s belief not need be doxastically justified since one may not believe it *for* the right reasons. Typically it is held that the so-called ‘basing relation’ marks the difference between a propositionally justified belief (if it is), and a doxastically justified belief, as I will explain below.

As many issues concerning the nature of propositional and doxastic justification should be left as open questions for the time being, such as for example, which of propositional or doxastic justification is more fundamental, in order to again avoid stipulating a substantive issue in the internalism / externalism debate. Also, the possibility that the fundamental natures of the two types of justification are quite different should be left open. For example, one might think of propositional justification in

¹⁰ See Feldman, R. and Conee, E., (2004) for a discussion of “well-foundedness” as a way of characterizing this distinction.

objective terms of logical or evidential relation, but might think of doxastic justification in subjective¹¹, responsiblist¹², or deontological terms¹³ (not that these evaluations are equivalent). Either way, when specifying and evaluating different forms of internalism/externalism, it is important to make explicit whether we are evaluating justification for *propositions* relative to a person, or *beliefs* that a person actually holds.

¹¹ Richard Foley is someone who thinks of *epistemic rationality* subjectively. While epistemic justification and epistemic rationality **may** be different properties, one could propose a theory of epistemic justification along the subjective lines that Foley does. To get an idea for what such a theory might look like, consider Foley's characteristic description:

It is egocentrically rational for you to believe a proposition only if you would think on deep reflection that believing it is part of what is involved in your having an accurate and comprehensive belief system. Your belief must be able to stand up to your own critical reflection. It must meet your own deep standards and in this sense be invulnerable to criticism. (Foley, 1993, p. 96)

Jose Zalabardo conceives of epistemic rationality in a similar way, though also incorporating elements of epistemic responsibility. Zalabardo defines epistemic rationality as follows:

ER2* If a subject has done her best by her lights to determine the truth value of a proposition *p*, then from the point of view of her conception of her epistemic situation, it is epistemically rational and responsible to believe that *p* just in case she believes that *p*. (Zalabardo, 2006, p. 24)

¹² Laurence Bonjour is a philosopher who closely links epistemic justification and epistemic responsibility. In fact, given what he says, he may even think of these notions as equivalent. Bonjour writes: "My contention here is that the idea of avoiding such irresponsibility [by accepting a belief in the absence of good reasons for thinking it true], of being epistemically responsible in one's believings, is the core of the notion of epistemic justification." (Bonjour, 1985, p. 8)

¹³ To say a belief is deontologically justified is not to say that the subject was obligated to believe that *p* (although that may also be the case), but that she was permitted to so believe, because so believing did not violate any epistemic duties and obligations. A duty is understood as an obligation or requirement. Richard Feldman suggests that an epistemic duty, specifically, is what is required for epistemic success. He claims that epistemic success is a matter of achieving epistemic goals. (Feldman, 2002, p. 376-7) What people take to be the epistemic goal varies, but standard candidates include the "truth-goal" (the goal of now maximizing truth and minimizing falsity in a large and important body of beliefs), or the "knowledge-goal", among others. So tying it altogether, Alston summarized statement of the deontological conception of justification is thus: "S is J-deont. in believing that *p* iff in believing that *p* is not violating any epistemic obligations." (Alston, 1989b, p. 87)

The Basing Relation:

The term 'basing relation' is a philosopher's term of art. However, our ordinary ways of speaking are sensitive to the distinction between basing a belief on a reason, and merely having the reason, but in no way 'connecting' the relevant belief with it. The epistemic basing relation is the relation that holds between a reason, or one's grounds, and one's belief when the belief is held *for* that reason. This contrasts with holding a belief and merely having a reason to believe it. Only in the former case is there a proper connection holding between a reason and a belief.

The basing relation is important to many areas in epistemology, particularly in regard to theories of epistemic justification that maintain that part of being justified is a question of believing for the right kind of reason, as opposed to merely having the reason in question. Advocates of doxastic justification will typically hold that the basing relation is necessary for such justification obtaining. For example, I might believe there is a conference this Friday. I might also have good reason to believe this having seen the date published on the normally highly reliable School of Advanced Studies website. Merely being in possession of this reason is not sufficient, however, to ensure that my belief is doxastically justified. I might believe the conference is on Friday because it is a full moon, and would have believed that the conference is on Friday whether or not I saw the website. It is only if I hold my belief because of, or in light of, my good reason - as we might say -- that the belief is, all else equal, doxastically justified.

The basing relation is necessary, but not sufficient, for epistemic doxastic justification.¹⁴ It is not sufficient because the basing relation simply concerns the manner

¹⁴ This of course presupposes the falsity of the most extreme versions of epistemic externalism, including what we might call 'mad-dog externalism', which is understood here as the view that the subject's

in which a belief ought to relate to a reason – it says nothing of what counts as a *good* reason, or for that matter, what a reason is. Furthermore, a properly based belief might still be unjustified, all things considered, if there remain undefeated defeaters. That is to say, if there are things (beliefs/facts/propositions, depending on one’s view) that undermine or rebut one’s justification that cannot themselves be undermined or rebutted by other considerations. In such a case one’s belief is not justified (and perhaps is unjustified) due to the presence of undefeated defeaters.¹⁵ Nevertheless, it is clear that the basing relation is at least necessary for doxastic epistemic justification.

There are two main construals of the basing relation. The first, and perhaps most common, are causal theories.¹⁶ Most generally, these maintain that a belief is based on a reason only if the reason suitably causes the belief.¹⁷ A causal theory of the basing relation is perhaps the default view since it requires the least amount of cognitive sophistication on the agent’s part, thus allowing it be applicable to the widest segment of epistemic beings. No higher-order concepts or cognitive abilities are presupposed. But due to two types of putative counter-examples to be discussed below, doxastic theories were developed as the rival account of the basing relation.

perspective is justificationaly irrelevant, and that factors external to it, e.g. the reliability of the process that gave rise to the belief in question, are necessary and sufficient for epistemic justification to obtain. On such accounts, reasons (traditionally construed) play no role in epistemology.

Reasons here are being understood broadly as those things which justify belief. Therefore, reasons might include belief states, as in the case of inferential justification, or for example, perceptual experiences, in the case of non-inferential justification. The claim that the basing relation is necessary for epistemic justification to obtain, therefore, should be palatable to Foundationalists, Coherentists, and Infinitists alike. More on how to best formulate these views below.

¹⁵ See 2.1. for a fuller discussion of defeaters.

¹⁶ See for example the account proposed in Moser (1989).

¹⁷ If one assumes that reasons are propositions, and propositions are abstract entities, it may seem odd to think of abstract entities as able to cause anything at all. Accordingly, throughout this thesis I will follow Robert Audi in interpreting “reason” in the locution, “a reason causes the belief that *p*”, as what he calls a “reason state”. A reason state is a causally efficacious mental state whose content is the proposition that functions as one’s reason to believe that *p*. So to say a reason *r* causes a belief that *p*, is shorthand for saying that it is one’s believing the reason *r* that is the cause of one’s belief that *p*. See Audi, (1993a), p. 234-5.

Doxastic theories require that the agent holds a meta-belief with the content such that the support relation is a good one that holds between a reason and belief. A *pure* doxastic theory, as I will call it, maintains that such a doxastic condition is *sufficient* for proper basing. A *hybrid* theory, as I will call it, maintains that such a doxastic condition is *necessary* for proper basing, in addition to a causal condition. Advocates of doxastic theories, such as Robert Audi, argue that such a meta-belief is required to ensure that belief is guided by the reason, not that the belief is a mere effect of the reason. (Audi 1993a) Presumably the reason why the content of the meta-belief must include the notion of a *good* support relation is that it is hard to see how one could rationally base a belief on something if one did not at least *take* the one to be good grounds for the other. If one believed the reason was a bad one, one should not hold the belief that would be based on it. If one had no meta-belief whatsoever about the support relation between ground and belief, it is hard to see a sense in which the belief could be *guided* by the reason, which is the chief concern of those who propose such theories.

As Keith Allen Korcz discusses in “Recent Work on the Basing Relation” (Korcz, 1997), a doxastic theory needs to overcome many potential problems. For example, he argues that if a reason is not another belief but merely a perceptual state, is it necessary that the subject have a meta-belief that a support relation obtains between the reason and the belief? One might think that this seems like an unrealistically stringent requirement. How could the subject believe, in a non-circular way, not only that his belief is based on a given reason, but also that the support is genuine? The worry seems in part to concern the possibility of a kind of non-inferential justification that would be needed to ensure that the belief is justified in a non-circular way that does not lead to vicious regress.

However, I do not see a problem here regarding the kind of justification that would be needed. My proposed elucidation of the Awareness Requirement below and the kind of non-inferential justification it contributes to can be applied here, *mutadis mutandis*. (See chapter 2)

Also, one might worry that many, if not most, individuals lack the epistemic concepts required to form the relevant meta-beliefs. Further, it is not obvious that a reason which grounds a belief will always be the same reason a subject would appeal to if asked why he believes something. While one might plausibly have privileged access to one's reasons, it is a much stronger thesis that one's reasons are always transparent to the subject, as well as claims that one is infallible about the basis of a given belief, which are claims the advocate of the doxastic theory might make.

That is not to say that these potential difficulties are necessarily insurmountable, but is a doxastic account of the basing relation motivated? This depends upon whether causal theories are really subject to the two standard objections commonly made against them. Accordingly, I shall examine the objections to a causal account of the basing relation and find them wanting. At best, the arguments establish that a *purely* causal account of the basing relation is incorrect. They fail to show that the basing relation is not, at least in part, a causal matter. Since these objections fail, the doxastic theory is unmotivated, at least via these considerations. If a theory that includes causal elements is broadly correct, the onus shifts back to the causal theorist to spell out the details of how such a theory can elucidate the basing relation.

A Causal Theory of the Basing Relation:

In a survey article, Korcz cites Paul Moser as an advocate of a causal theory of the basing relation. (Korcz, 1997) In fact, given what Moser says, this seems to be an incorrect classification of his theory, as I will explain below. According to Korcz, Moser holds that:

S's believing or assenting to p is based on his justifying propositional reason q = df. S's believing or assenting to p is *causally sustained* in a non-deviant manner by his believing or assenting to q , and by his associating p to q . (Korcz, 1997, p. 2) (emphasis added)

Moser's account is construed in terms of what it takes for one belief to be based on another. While he does distinguish between non-propositional reasons and propositional reasons, for simplicity and clarity of exposition I shall focus only propositional reasons throughout the rest of this discussion.

Korcz takes issue with the second conjunct in Moser's account, i.e. that the subject must in some way associate belief p with its supporting reason q , or at least be disposed to. This requirement reveals that Moser does not hold that causation alone is sufficient for basing. Accordingly, it is at best misleading and at worst false to claim that he is an advocate of a causal theory of the basing relation. I will call a causal theory of the basing relation one that holds that a causal relation is *sufficient* for basing. Given this distinction, it seems Moser is an advocate of what I will call a *hybrid-causal* view of the basing relation, e.g. this is the view that causation alone is insufficient for basing, but necessarily requires further conditions to be met, for example an appreciation of the evidential connection between a belief and the grounds it is based on. Of course it is an open question of what such 'appreciation' consists.

Korcz objects to Moser's view on the grounds that someone might come to believe something on the basis of a reason without realizing that one supported the other. For example, Korcz argues that it seems possible that one could genuinely base a belief on subliminal advertising. (Korcz, 1997, p. 2) In virtue of the reason being subliminal, one would never be able to appreciate the support relation, and therefore could not associate the belief with a reason. Still, the belief is no less based on a reason. Or so Korcz argues. The opponent of the causal view can concede that it is possible that a belief could be *caused* by a subliminal reason. The question is if this is enough for it be *based* on that reason in a way that is epistemically relevant, e.g. could such a subliminal reason contribute towards the *justification* of the belief that is caused by it? I maintain that it cannot, because such a belief would be subject to what I will call the Subject's Perspective Objection (SPO) (see chapter 2 for elaboration).

However, even if Korcz's criticism is sound, it counts only against Moser's formulation of the basing relation and others that include an 'association' requirement. That says nothing against causal and hybrid-causal theories of the basing relation in general. In what follows, for brevity I will focus on pure, rather than hybrid, accounts of the basing relation unless I indicate otherwise. If causation is not necessary for basing to obtain, this will obviously count against both pure and hybrid causal accounts of the basing relation. Accordingly, I will now turn to what are taken to be the two standard objections to theories that include causation as an essential element of basing.

Deviant Causal Chains and the Case of the Gypsy-Lawyer:

The first counterexample offered is called the Problem of Deviant Causal Chains. It is essential to a causal account of the basing relation that the belief be causally sustained by the reason. The problem highlights that not just *any* kind of causal chain will be sufficient for a belief to be based on a reason. Korcz directs the reader to an example Alvin Plantinga gives of a deviant causal chain. In a footnote, Plantinga writes,

Suddenly seeing Sylvia, I form the belief that I see her; as a result I become rattled and drop my cup of tea, scalding my leg. I then form the belief that my leg hurts; but though the former belief is a (part) cause of the latter, it is not the case that I accept the latter on the evidential basis of the former. (Plantinga, 1993, p. 69)

Note of course that advocates of a causal theory never held (or should not have held) that just *any* causal relation between beliefs is sufficient for one to be based on the other. Recall for instance that Moser explicitly works into his basing condition that one belief must causally sustain the other in a *non-deviant* manner. How a causal theorist might respond to this challenge will be revisited later.

The second type of counter-example to causal theories of the basing relation have been called “Gypsy-Lawyer” counterexamples.¹⁸ Where the problem of deviant chains tries to establish that the causal condition is insufficient for basing, Gypsy-Lawyer cases purport to show that a causal condition is not even necessary for the basing relation to obtain.¹⁹ There have been two broad types of responses to Lehrer’s example: to either

¹⁸ Keith Lehrer first proposed the Gypsy-Lawyer counterexample in Lehrer (1971). This counterexample runs throughout his work, however. See for example Lehrer (1974), Lehrer (1990) and Lehrer (2000).

¹⁹ While Lehrer himself presents his case as a counterexample to whether a belief based on a reason needs to be *explained* by that reason, all of the authors I cite here who endorse the counterexample understand the kind of explanation in question as *causal*. So regardless of what Lehrer himself originally intended, gypsy-lawyer style counterexamples have been standardly taken to count against the necessity of causation to the epistemic basing relation.

accept it as genuine, and to modify one's theory to respect it²⁰; or to reject it out of hand, without giving a clear account of what is wrong with the counterexample. For example, in response to the case, Alvin Goldman, writes: "I find this example unconvincing. To the extent that I clearly imagine that the lawyer fixes his belief solely as a result of the cards, it seems intuitively wrong to say that he *knows* — or has a *justified belief* — that his client is innocent." (Goldman 1979, p. 22, n. 8)

I will argue below that so-called Gypsy-Lawyer counterexamples do not count against theories of the basing relation that include a causal element, since such counterexamples are incoherent as presented. I will be explicit about both what is wrong with the examples, as well as why people have been mistakenly taken in by them.

Since the example is quite complicated and much depends on the details, I will quote the key elements verbatim from Keith Lehrer's original presentation of the counterexample. Lehrer explains the structure of his example as follows:

In my example a man comes to believe something and continues to believe it because of groundless superstition. As a result of having the belief, he uncovers reasons for the belief that give him knowledge. But these reasons do not potentially explain his belief, because he would not hold the belief for those reasons if he were to become doubtful of his superstitious reasons for belief. (Lehrer, 1971, p. 311)

Lehrer presents the case of a gypsy lawyer who is defending his client on charges of eight brutal murders. There is overwhelming evidence against the client for the first seven murders and all, including the lawyer, believe the accused is guilty as charged of committing all eight. The lawyer, seeking guidance, turns to his tarot cards and will believe what they tell him unquestioningly. The cards, unbeknownst to the lawyer, are in

²⁰ For example, Swain (1979), (1981), (1985); Tolliver (1982); Korcz (2000) have all accepted the Gypsy-Lawyer counterexamples as genuine.

fact utterly unreliable. The cards tell the lawyer that his client is innocent of the eighth murder.

Realizing he needs something less controversial to offer in court than the say-so of the tarot cards, he re-examines the evidence and finds a complicated line of reasoning which in fact deductively proves his client is innocent of the eighth crime. Lehrer claims that this line of reasoning gives the lawyer knowledge of his client's innocence, and in fact the lawyer claims that it is this reasoning that gives him knowledge. (Lehrer, p. 312) Lehrer goes on to say that no one else is convinced by the sound but highly complicated reasoning because it is so complicated and subtle, combined with powerful emotional factors that surround the case. In fact,

and this is the crucial point, if it were not for his unshakeable faith in the cards, the lawyer himself would be swayed by those emotional factors and would doubt that his client was innocent of that eighth murder. It is only because of his faith in the cards that the reasoning gives him knowledge. (Ibid.)

Lehrer concludes that,

Therefore, the reasoning that gives him knowledge does not even potentially explain his belief that his client is innocent of the eighth crime. It in no way supports or reinforces his belief, but it does give him knowledge. (Ibid.)

From examples like this, the opponent of causal analyses of the basing relation concludes that what causes the gypsy's belief that his client is innocent of the crime is the cards, but the justificatory basis of his belief is the complicated line of reasoning. In short, they hold that a reason causing a belief is not necessary for the belief to be based on that reason.

Discussion:

In this section I will argue that the Problem of Deviant Causal Chains does not count against causal theories of the basing relation, and neither do the Gypsy-Lawyer counterexamples since they are incoherent as presented.

Recall that the Problem of Deviant Causal Chains attempts to undermine the sufficiency of a causal theory of the basing relation. As I alluded to in the exegesis above, no one thought (or ought to think) that just *any* kind of causal chain is sufficient for the basing relation to obtain. As Moser stresses, the causal chain must be “non-deviant”. What non-deviance consists of, however, is obviously a hugely difficult and possibly unsolvable problem.²¹ I will attempt no analysis of non-deviant causation here. But from a lack of analysis of non-deviance (or at least unanimous agreement of what non-deviance consists of), it does not follow that a causal condition of the basing relation is insufficient. Compare this with analyses of what it is for one event to cause another. No successful analysis of causation has been provided (or none that is widely agreed upon). From this it does not follow that events do not cause one another, or that one event is never sufficient to cause another, only that we have no reductive analysis of the relation.

The lesson to take from the Problem of Deviant Causal Chains, I suggest, points to a general problem about the analysis of causation, not specifically against causal accounts of the basing relation. At most, what the problem of deviant causal chains shows is *purely* causal theories of the basing relation are inadequate. I submit that the reason that Moser and others introduce a condition that requires that the subject

²¹ That is, the problem may have no solution, if by ‘solution’ one is seeking a reductive analysis in terms of necessary and sufficient conditions of what it is for one thing to cause another.

appreciate the relation between ground and belief (e.g. by having a further meta-belief, or through conscious awareness) is precisely to rule out cases of deviant causation. But that is not to say that causation does not play an essential role in proper basing.

As for Gypsy-Lawyer counterexamples, I find they are incoherent as presented, and hence do not show that a causal element is not necessary for the basing relation to obtain. The conjunction of the following three propositions cannot be true:

Recognition: “The lawyer *recognizes* that the complicated line of reasoning shows that his client is innocent, but the complicated line of reasoning cannot cause the lawyer to believe that his client is innocent due to the emotional factors surrounding the case.” (Korcz, 1997, p. 3) (emphasis added) Further, the line of reasoning is the sole basis of the justification for his belief that his client is innocent.

Cards: The lawyer’s belief is causally explained by the tarot cards alone

Rational: The lawyer is rational (narrowly construed)²²

I contend that this triad is inconsistent, and thus consistency requires denying one or more of the propositions. That is, if *Recognition* is true, either *Cards* or *Rational* has to be false, and similarly for other possible combinations. However, in that case we would not have a Gypsy-Lawyer counterexample. It seems that for such a counterexample to go through, the three propositions must be compossible. But they are not.

Lehrer does not explicitly state *Rational*, but since he does not say anything about the lawyer’s rationality one way or the other, one can only charitably assume the default

²² Following Thomas Scanlon, here I will construe rationality narrowly. That is to say, a person is rational when their attitudes conform to their own judgments about what reasons are present. So, for example, a person is irrational in this sense if he continues to believe something even though he judges there to be good reason to reject it. See Scanlon (2000), p.25. In such a case the person believes “contrary to the person’s own reason”. (Ibid.) In such a case the person recognizes something as a reason but fails to be affected by it in one of the relevant ways, e.g. adopting or ceasing to have the attitude. In construing rationality in this narrow sense I am not claiming that this is the *only* sense of rationality there is, I am just committing myself to the claim that there is at *least* this sense of rationality.

view that the lawyer is rational (in at least this narrow sense). This is not undermined by the fact that the lawyer is a devout believer in tarot cards, since this narrow sense of rationality only requires consistency between attitudes the subject has and their judgements about what reasons are present. A subject can be rational in this narrow sense while being manifestly irrational in other senses. This narrow sense of rationality says nothing about what judgements the subject *should* make about the reasons, or what good reasons consist of. As such, it is a fairly undemanding condition, and so can be unproblematically assumed in this context.

So if it is the case that the lawyer is rational (narrowly construed), and the lawyer's belief is genuinely caused by his tarot card belief, *Recognition* must be false, since it is hard to see how he can recognize or otherwise appreciate that a consideration counts in favour of something, and yet rationally fail to believe it. If *Recognition* and *Rational* are both true, then *Cards* is strictly false, but perhaps he believes on the basis of both *Recognition* and *Cards** (*Cards** might be something like: the lawyer's belief is at least *partly* causally explained by the tarot cards.) But again, this would not be a Gypsy-Lawyer counterexample since such examples purport to show that causation is not necessary *whatsoever* to basing. If *Rational* is false, then it becomes very difficult to judge what belief justifies what since the agent's beliefs are not sensitive to considerations that he takes to count for or against them.

The reason to think *Recognition* is false is because it is unclear what *recognizing* that a reason shows a conclusion is true amounts to, short of believing the conclusion on the basis of the reason (or, if the recognition is non-doxastic, it will necessarily dispose a rational agent to believe the conclusion on the basis of the reason). It does not seem

rationally possible to believe (or otherwise recognize) that something is sufficient evidence for something else, believe that evidence, and then not believe (or be disposed to believe) the supported belief.²³

To think otherwise leads to paradoxical results, such as those that arise in Moore's Paradox, e.g. "It's raining but I don't believe it." If, by asserting that it is raining, one is representing oneself as believing that it is raining, it is quite odd in the next breath to deny that one believes it. If the lawyer, in fact genuinely recognizes that the reasoning proves the innocence of his client (rather than merely saying he recognizes it, without doing so), he *ipso facto* believes his client is innocent. In that case it is reasonable to assume he believes on the basis of the complicated reasoning. That is, it is his belief about the complicated line of reasoning that at least partly causes his belief that his client is innocent. Otherwise, it is difficult to see how he could still be a rational believer.

If *Cards* and *Rational* are true and *Recognition* is false, how can we explain why those like Lehrer mistakenly conclude that the lawyer does base his belief on the complicated line of reasoning, even though it does not causally explain why he holds the belief? I think a reason why they might be misled is because they conflate basing a belief on a reason, and *potentially* basing a belief on a reason. This is analogous to the distinction between a belief being justified (doxastic justification) and merely its propositional content being justified for a person (propositional justification). If a belief is propositionally justified, then it is a candidate for being doxastically justified in the sense that the right logical relations hold between the relevant propositional contents.

²³ Cf. Shah (2003) on the 'transparency of truth'; that is, the phenomenon of why when deciding to believe that *p*, one only has to determine whether *p* is true – there is no gap between them. Instead, the former question collapses into the latter.

When a belief is doxastically justified, however, the subject has made the relevant connections between the propositions, i.e. the belief is *based* on its grounds.

Similarly, in the case of the complicated line of reasoning, the lawyer's belief can potentially be based on the reasoning since it is information that he has available to him. Perhaps then he is propositionally justified in believing his client is innocent of the eighth crime. It is only when he accesses this information and comes to accept it by believing it on that basis, however, that his belief is doxastically justified.

Lehrer and Korcz's discussion of the case of the Gypsy-Lawyer gets further muddled by talk of epistemic notions like justification and knowledge. Some discussion revolves around whether the lawyer is justified on the basis of the reasoning or the cards, or if he knows his client is innocent. This talk of justification and knowledge are potentially misleading. It is imperative to keep clear the distinction between propositional and doxastic justification since no one disputes that *if* the lawyer bases his belief on the cards, he is unwarranted (i.e. due to their unreliability, he lacks Warrant, understood here as that which is required to convert true belief into knowledge), even though the cards *may* provide propositional justification. Likewise, no one should dispute that if the lawyer bases his belief on the reasoning, his belief will enjoy doxastic justification (and possibly Warrant, assuming, again, that there is such a thing). The real question is whether someone can base a belief on another if the latter does not, at least in part, causally sustain the former. There is no problem about basing belief on bad reasons. The question is about the basing relation *simpliciter*.

Someone might object with the following kind of putative counterexample. Suppose a subject believes that *p* for a perfectly good reason *q*. The subject believes that

p on the basis of reason q . Sometime later, the subject acquires the additional reason r . Intuitively, r is a reason for p , but r did not cause p , since p was already held and the belief was not modified in any way when the subject acquired the additional reason. If this is correct, does this show that reason r has not caused anything in relation to the belief that p ? And if we can say that the subject has acquired this reason, and his belief is based on it since it justifies his belief, does this not show that causation is not necessary for the basing relation after all?

In reply to this possible case, it is true that the subject believes that p on the basis of reason q , and then acquires the additional reason r . It is also correct that intuitively r is a reason for p , but r did not cause p , at least in the sense of *initially* causing p , since p was already held. Despite all this, r *causally sustains* p . It is just that the subject's belief is justificationally overdetermined, so to speak, in the sense that the subject has two or more reasons that are individually sufficient for the justification of what is believed. But r causally sustains p nevertheless. Plausibly, if the subject lost his reason q (he forgot it, say), he would still justifiably believe that p on the basis of r . And if the subject lost r (he forgot it, say), if rational, he would drop his belief that p . A subject's reason does not need to be the initial cause of what is believed for proper basing to take place, but a causal connection of some sort necessary, e.g. a reason that causally sustains the belief. The epistemic basing relation is, at least in part, causal.

By analogy, consider the fact that I am alive. The (a) reason for my existence is my parents. An additional reason for my continued existence is the proper functioning of my lungs. Intuitively, the proper functioning of my lungs is a reason for my continued existence even though their operation was not the initial cause of me, since presumably I

was already alive before I developed lungs in utero. Nevertheless, my lungs causally sustain my existence. If I lost my parents, I would still be alive, but the reason would be (in part) due to the proper functioning of my lungs. In turn, if my lungs failed, I would die (assuming no one intervened), which supports the claim that they (at least in part) causally sustain my life. The relation here is at least in part causal, and this is not undermined by the fact that there can be genuine reasons for my continued existence that enter into the explanation downstream from my genesis. As we saw, similarly in the case of reasons for belief.

In short, if the lawyer's belief that his client is innocent is based on the complicated line of reasoning as *Recognition* contends and he is rational, the relation between his belief about the complicated line of reasoning and his belief about his client's innocence must be in part causal. In that case the belief about his client is caused by both the reasoning and the cards, or if the causal chain leading from the cards to the belief is deviant, then on the complicated reasoning alone. Either way, the lawyer's belief his client is innocent of the eighth crime is (at least in part) causally sustained, and therefore a candidate for being based, on his reason(s).

This discussion has primarily considered objections to a causal theory of the epistemic basing relation. That theory tends to be taken as the default view because it is the least intellectually demanding account of the basing relation, requiring no higher-order cognitive abilities such as the ability to form meta-beliefs about one's grounds. As such, the theory is aimed at the widest segment of cognizers. Also, and perhaps most fundamentally, if something is a reason for someone's belief, it just seems that the most obvious suggestion is that there is some kind of causal link between them. We are

assigning the reason an explanatory role, and causation seems to be the most obvious candidate to, at least in part, fill that role.²⁴ These considerations are not conclusive, but they do justify a (defeasible) presumption in favour of a broadly causal theory of proper basing. Therefore, the aim was to disarm objections to the default view, leaving more complicated rival theories unmotivated.

To that end, it was shown that the deviant causal chains problem is no objection in itself to the sufficiency of a causal account of the basing relation. However, for reasons that I will explain in the next chapter to do with the subject's perspective, a purely causal theory is the incorrect one. Such considerations count in favour of a hybrid-causal theory of the basing relation that requires the subject's conscious awareness of the relation between his grounds and his belief. This will not be a *doxastic* theory of the basing relation, since I will argue that there are forms of non-doxastic awareness that can play the same role a meta-belief is supposed to in doxastic theories. Such an approach will share the benefits of a doxastic approach (like Audi's), without sharing its pitfalls (such as vicious regresses, etc.) Deviant causal chains merely highlight a concern relevant to causation generally, not a problem unique to the basing relation. As in the case of causation generally, deviant causation counts not against causation itself, but only against our current analyses of it, or perhaps against the possibility of giving a reductive analysis of causation at all.

It was then shown that Gypsy-Lawyer counterexamples do not challenge the necessity of a causal element of the basing relation because the counter-examples are incoherent as presented. Assuming the lawyer is rational (in the narrow sense), he cannot

²⁴ Cf. Davidson (1963).

fail to base his belief on a reason he accepts as sufficient for its truth. The relation here is, at least in part, causal.

With objections to the basing relation that incorporate causation dispensed with, the task of advocating such theories is to spell out the details of such an account, keeping in mind the issue of deviant causation. The initial minimal necessary assumption, however, is that for a belief p to be based on a reason r , r must non-deviantly causally sustain p .

With a discussion of the basing relation in place, we are now in a position to explore ways of construing the internalism/externalism distinction.

1.4. Some Ways of Construing the Internalism/Externalism Distinction:

Typically, epistemic internalism/externalism is specified relative to the *grounds* or to the *justificatory status* of a belief that the subject holds.²⁵ In the former case, one might insist that for a belief to be justified one must have awareness of the mere existence of one's grounds. Following something of a convention established by Jim Pryor, I will call this *Simple Internalism*. Jim Pryor defines Simple Internalism as the view whereby "whether one is justified in believing *p* supervenes on facts which one is in a position to know about by reflection alone." (Pryor, 2001, p. 104) On this view, in order for a fact to contribute to justification, its presence must be accessible to the agent.²⁶

In the latter case, one might insist that for a belief to be justified one must be aware of the fact *that* one is justified. Again following Pryor, I will call this *Access Internalism*. An access internalist maintains that "one always has 'special access' to one's justificatory status." (Pryor, 2001, p. 105) "Special access" here is understood as *a priori* ways of coming to know. So, unlike Simple Internalism that only requires access to the presence of one's grounds, or *Inferential Internalism* that only requires access to reasons to think one's non-basic beliefs are well supported, Access Internalism makes the strong requirement that *all* of one's beliefs, basic and non-basic alike, are such that not only are one's grounds accessible, but also the fact that the grounds are adequate.²⁷

²⁵ See Alston (1989e) for a survey of ways of construing the internalism / externalism distinction.

²⁶ William Alston defends a paradigm case of a Simple Internalism. In "An Internalist Externalism" he defends the view that in order to be justified, there must be a kind of cognitive access to the ground of the belief. (Alston, 1989c, p. 237-239) On the other hand, he argues that the adequacy of the ground, i.e. its truth-conduciveness, is an external fact that the subject need not be able to access. (Ibid., p. 239-244) As Alston summarizes in the conclusion of that paper, "I have given reasons for placing a (rather weak) AI [access internalist] constraint on something's being a ground that could justify a belief, but I have resisted attempts to put any internalist constraint on what constitutes the adequacy of the ground." (Ibid., p. 244) Hence an internalist externalism.

²⁷ For a defense of inferential internalism, see Fumerton (1995), ch. 3.

That is, on one understanding of Access Internalism, for any justified belief *p*, the subject must be able to tell upon reflection alone that the belief is justified. An advocate of this kind of Access Internalism might embrace a “JJ Principle”, which says that if one has a justified belief then one must have justification to believe that one is justified, which is similar to the “KK Principle”, but cast in terms of justification instead of knowledge. (Ibid.)²⁸ For example, take this characteristic expression of access internalism from Chisholm about what the subject can determine about his own doxastic state *a priori*:

Merely by reflecting upon his own conscious state, he can formulate a set of epistemic principles that will enable him to find out, with respect to any possible belief that he has, whether he is *justified* in having that belief. The epistemic principles that he formulates are principles that one may come upon and apply merely by sitting in one’s armchair, so to speak, and without calling for any outside assistance. In a word, one need consider only one’s own state of mind. (Chisholm, 1989, p. 76)

When considering the kind of “special access” that might be required here for an internalist theory of justification, it is worth noting that it *need* not be construed so narrowly as to include *only introspection* as Chisholm suggests. I think that the internalist’s insistence on introspection has resulted in focusing on cases of memory and/or inferential justification. It is presumably myopia that has led to overlooking central cases of justified believing such as having beliefs about what one is presently experiencing through sense perception. In that case one does not *introspect* the world around one but rather one *consciously experiences* it. What is notable here is that the internalist should maintain that for every occurrent justified perceptual belief about one’s external environment at time *t*, there must be at least one conscious sense experience at *t*.

²⁸ For classic expressions of access internalism, see Bonjour (1985) and Chisholm (1989).

That is not to say that there must be a one-to-one correspondence of occurrent beliefs to perceptual states – there may be an indefinitely many beliefs that can be justified by a given experience. What is common to perception and introspection that the internalist is moved by is *awareness*, which I will capture in the Awareness Requirement on justification below (see Chapter 2).

In sum it is not that there is anything particularly cogent for the internalist about introspection *per se*. What is justificationaly important about introspection is that it is our way of becoming *consciously aware* of the contents of our minds, just as perception is our way of becoming *consciously aware* of the world.

Besides grounds and adequacy of those grounds, we can now see that there is a third element that is relevant to doxastic justification on which internalists/externalists can take a stand – that is, the basing relation. (Swinburne, 2001, p. 11) Is the fact that a belief is based on a reason something that the subject must be aware of, as a doxastic theory typically holds, or be “internal” in some other sense, e.g. must the subject be ‘non-doxastically’ aware of an evidential fit between a belief and the reason it is based on? Or can the mere fact that proper basing occurs be an external fact as advocates of *purely* causal accounts of the basing relation suggest?²⁹ I will argue that it is a necessary condition of a belief being based on a ground in a way that can yield doxastic epistemic justification that the subject meets various awareness requirements, although I maintain that there are non-doxastic forms of awareness, viz., that there are forms of conscious awareness without belief. (See Chapter 2)

²⁹ Recall that this is the view that causation alone is sufficient for basing, and therefore further conditions need not be met, such as an appreciation of the evidential connection between a belief and the grounds it is based on.

It should be clear, however, that the stance one takes on the basing relation will shape how internalist/externalist one's theory of epistemic justification will be. In principle, a theory of justification could blend internalism/externalism in regard to any of the three elements of grounds, adequacy of the grounds, or the basing relation.

In summary, so far by epistemic internalism we have meant that what is justificationaly relevant must be, or easily can be, the object of conscious awareness. But besides emphasizing the justificatory role of the conscious awareness, the other main way of construing the 'internal' of epistemic internalism is the 'Mentalist's' construal of the 'internal' as supervening upon one's intrinsic physical states. In chapter 2 I will offer a full discussion, and ultimately reject, Mentalism as its proponents formulate it.

But even among those who stress the epistemic significance of conscious awareness, we have seen diverging answers as to what kinds of things are taken to be justificationaly relevant by different theorists:

- i) Grounds: *Simple Internalism* is being understood as the view that whether one is justified in believing that *p* supervenes on facts which one is in a position to be consciously aware of. In this view, in order for a fact to contribute to justification, its presence must be accessible to the agent.
- ii) Adequacy of Grounds: *Access Internalism*, by contrast, maintains that one always has 'special access' to one's justificatory status. So unlike simple internalism which only requires access to the mere presence of one's grounds, access internalism makes the strong requirement that *all* of one's beliefs, basic and non-basic alike, are such that not only are one's grounds accessible, but also that the grounds are adequate. That is, for any

justified belief p , the subject must be able to tell upon reflection alone that the belief is justified.

- iii) Adequacy of Grounds, Redux: *Inferential internalism* was introduced as a sub-species of Access Internalism in that it only requires access to reasons to think one's *non-basic* beliefs are well supported. In section 1.5. below I will elaborate on the distinction between basic and non-basic beliefs and defend the view that many of our basic beliefs are justified.
- iv) The Basing Relation: Is the fact that a belief is evidentially related to a reason something that the subject must be aware of if one is to be based on the other, as for example a doxastic theory of the basing relation typically holds? A theory that answers this question affirmatively as a necessary condition of epistemic justification obtaining is 'internalist' in another distinct sense.

The basing relation is not superfluous since it distinguishes two different *kinds* of epistemic justification, even among theories that embrace awareness requirements. One might have thought that if justification depends on being based on a ground, then given an awareness requirement, how can one know that one has a justified belief without knowing what the ground is? But this is to focus only on *doxastic* justification. *Propositional* justification obtains merely if the subject has reason to believe the proposition in question, independently of the question of whether the subject holds the relevant belief, or whether the subject bases the belief he does hold on the relevant ground.

Having explored different ways of construing the internalism/externalism distinction, and before defending my preferred version of internalism in chapter 2, I will conclude my preliminary exploration of the basic features of epistemic justification by addressing the question of the *structure* of epistemic justification.

1.5. The Structure of Justification and the Existence of Immediate Justification:

This section will explore what immediate justification is, as well as if we have any. After concluding that putative examples of immediate justification *prima facie* support the existence of immediate justification, I will consider the primary argument against it. Discovering what is wrong with the argument against immediate justification draws attention to important constraints on an adequate theory of epistemic justification (including an *internalist* theory of epistemic justification).

On Immediacy:

With a rough understanding of what justification is, by having considered among other things its four basic features as well as doxastic v. propositional justification and the basing relation, it is necessary to turn to the question of whether it can be had immediately. Pryor defines immediate justification as follows: “When your justification to believe *p* does *not* come from your justification to believe other propositions, I will call it *immediate*.” (Pryor, 2005, p. 183) By contrast, “When your justification to believe *p* comes in part from your having justification to believe other, supporting propositions, I will say that those latter propositions *mediate* your justification to believe *p*.” (Ibid.) This distinction is also known as the one between basic and non-basic beliefs as well as the one between non-inferential and inferential beliefs. As these labels are equivalent to the immediate/mediate distinction, they could be used interchangeably.

However, I think that the label “inferential/non-inferential” justification should be avoided since it is apt to mislead. By putting the distinction in terms of *inference*, however, one might wrongly think that the issue is the *psychological* process of how the

belief was formed. Usually inference is considered to be the process of passing from belief in the premises of an argument to belief in another premise, or a transition from premise to conclusion. The mediate/immediate justification distinction, on the other hand, concerns the kind of *justificatory* support relations that obtain between beliefs and different kinds of grounds, not the psychological process of how the subject formed the belief in question. To further clarify what is what is at issue, Pryor gives as an example of a proposition that is justified immediately, “(Gauge) The gas gauge reads ‘E’”, based upon one’s experience of looking at one’s petrol gauge. (Pryor, 2005, p. 182) A mediately justified proposition, on the other hand, is “(Gas) Your car is out of gas”, which gets its justification from Gauge and other justified propositions one holds. (Ibid.)

Pryor’s focus on one’s justification to believe rather than on whether one is actually justified is brought out in a comment he makes about the case of (Gas) and (Gauge) above. He says, “it is not important for our purposes whether you *actually do* believe (Gauge) or (Gas). Given your evidence, you *ought* to believe them.” (Pryor, 2005, p. 182) But perhaps this latter claim is too strong in emphasizing that the subject *ought* to believe the propositions. Nevertheless, surely one would be justified in believing them, were one to form them on one’s available evidence. One need not make the stronger claim that Pryor does to make this claim about what one is justified in believing. So whether or not Pryor is correct about what the subject ought to believe in this case, the emphasis is on having reasons to believe the conclusion, whether or not the subject actually holds the belief. What is crucial here is the idea of justificational immediacy.

It is important to note what kind of independence immediate justification is supposed to embody. As Pryor understands immediate justification, it is a property of a belief that is *justificationally* independent of other beliefs, that is, a belief is immediately justified if it does not derive its *justification* from standing in relation to other beliefs. Although it is important to distinguish, as I do below, between justification coming from the justification of other beliefs and justification coming from beliefs whether they are justified or not (if such a case is possible). However, it is perfectly compatible with a belief being justificationally basic that it is necessary one have other beliefs or concepts in order to *form* the basic belief in question. (Pryor, 2005, p. 183)

This point parallels one made in debates about the *a priori*. There the idea is that a belief is *a priori* if its *justification* is independent of sense experience, even if the concepts and other beliefs required to form the belief are learned through sense experience. For example, one may learn the concepts of ‘bachelor’, ‘marriage’ and ‘man’ through experience, but one’s belief that bachelors are unmarried men is *justified a priori* (presumably it is one’s conceptual understanding, or some other non-empirical source, that is the source of one’s justification).

A final point to help clarify a possible objection about how to understand immediate justification: in the case of mediately justified beliefs, the grounds are other beliefs. In the case of immediately justified beliefs, the grounds are *usually* something other than beliefs, e.g. experiences of certain kinds. Exceptions to this are beliefs about our own beliefs. Take the belief that *Freddy is Canadian*.³⁰ I can then reflect on my belief and form the further belief, that *I believe that Freddy is Canadian*. In this case it is a belief with another belief embedded in it, namely the belief that *Freddy is Canadian*. It

³⁰ Italicized words will be used to designate the propositional content expressed by those words.

does not seem possible for me to think *I believe that I believe that Freddy is Canadian* without believing that *Freddy is Canadian*. That is, the latter belief is required in order to form the former.

Assuming that my belief that *I believe that Freddy is Canadian* is justified, one might ask, is it justified mediately or immediately? Given the clarifications I made above, it should be clear that it is *immediately* justified. That is, my *justification* for my belief that *I believe that Freddy is Canadian* does **not** derive from my belief that *Freddy is Canadian*. I might be entirely unjustified in believing that Freddy is Canadian. I might have no reason to think that he is Canadian, compounded by the fact that I know he was born and raised in Yorkshire, speaks with a thick Yorkshire accent, has no Canadian relatives that I have ever met, and as far as I know has never so much as even visited Canada.

Of course all of this is not incompatible with his actually *being* Canadian (perhaps he married a Canadian and became a naturalized citizen). But if I have no positive reason to believe that he is Canadian and have strong evidence that he is English, my belief that he is Canadian is quite unjustified. But this need not stop me from believing it (if only it did!) My belief that *I believe that Freddy is Canadian* could be justified, even if my belief that he is Canadian is not. This shows that the justification for my belief about my belief is not *derived* from any other beliefs, since here a belief that is needed in order to form this belief is not itself justified, and so has no justification to give. Therefore, beliefs about beliefs can be immediately justified.

What this shows is that the category of immediate justification is heterogeneous. One might believe that there is a cat in front of one since one at least seems to see a cat.

This is a case of immediate justification, since what justifies one's belief is not another belief, but rather a perceptual experience of a certain kind. But what justifies one's beliefs about one's beliefs, whatever that is, is not like the perceptual belief case. So, somewhat oddly, the category of immediate justification must be negatively defined.

It is important to note that what is at issue here for justification for beliefs about beliefs is not that they are psychologically spontaneous. When entertained, we tend to assent to their truth *immediately*, but only in the sense of spontaneously, without reflection. The issue here concerns the structure and kinds of grounds that comprise the *justificatory* support relations, not the psychological process whereby individual beliefs are formed, e.g. if they are formed spontaneously, after conscious inference, etc.

Is There Immediate Justification?:

As I understand immediate justification, a belief is mediately justified if its justification depends on having at least one other belief, whereas a belief is immediately justified if its justification does not derive from standing in relation to any other belief. The traditional argument for this position has been an argument from elimination. In such an argument, advocates of Foundationalism argue that all other possibilities lead to sceptical conclusions or objectionably circular reasoning, and that only their view gives the correct account of the structure of justification. The traditional argument employed to this effect is an epistemic regress argument known as Agrippa's Regress. The regress begins by noting that if some justification is inferential or mediate (which it surely is), then the choices one faces are as follows:

1) All justification is inferential and an infinite regress looms

Advocates of this structural position can either argue for the sceptical conclusion that the regress is vicious, or else they argue that the regress can be virtuous (a position known as Infinitism).³¹

2) All justification is inferential but eventually loops back on to itself in a way that yields justification (Coherentism).

3) The regress terminates in some beliefs that are properly basic/or immediately justified, and all non-basic beliefs ultimately derive their justification from standing in relation to the basic beliefs (Foundationalism).

As Pryor notes, the weakness of this type of argument is that it relies on establishing that the rejected alternatives really are unworkable. (Pryor, 2005, p. 184) Pryor presents an unusually simple argument for Foundationalism, namely, an argument from examples.

Pryor merely cites examples in which it clearly seems that they are instances of beliefs that are immediately justified, and argues that it is implausible to suggest that they are justified by other propositions one has justification for believing. Take these examples from Pryor: “Suppose I feel tired, or have a headache. I am justified in believing I feel those ways. And there does not seem to be any other propositions that mediate my justification for believing it. What would the other propositions be?” (Pryor, 2005, p. 184) What seems to justify one’s belief about one’s own headache when one has a headache, for example, is one’s experience of one’s aching head. It is implausible to suppose that to be justified in believing that one has a headache one must also have beliefs about the features in question, e.g. one does not infer from a belief about the

³¹ For a defense of Infinitism, see Klein (2005).

experienced pain of one's aching head and a believed principle that 'if one has an aching head, then one has a headache', that one has a headache! While one may often *also* have beliefs about the feature in question, e.g. about the pain in one's head, it is implausible to suppose that these beliefs *mediate* one's justification for one's belief that one has a headache. If we stipulate that one lacks a belief about the pain in one's head, and in the mere presence of one's experience of one's aching head, one is justified in believing that one has a headache.

Pryor cites other examples that include justified beliefs about why one performs simple bodily movements and justified beliefs about certain conscious mental acts one undergoes, such as imagining. Pryor suggests, and I think rightly, that since these seem like overwhelmingly intuitive cases of immediate justification, the burden of proof is on those who oppose immediate justification to show how these kinds of beliefs are not really immediately justified after all. The best argument against the possibility of immediate justification might be called the Master Argument for Coherentism.

Searching for the "Master Argument" for Coherentism:

I will follow standard usage and refer to the view that all justification is mediate (and infinite regresses cannot yield justification) as Coherentism. In addition to these explicit features, the key to characterizing this position as a version of Coherentism is to understand the *mediation* relation between beliefs as a matter of coherence or fit. After some refinement, Pryor presents what he calls the Master Argument for Coherentism:

The Content Requirement (Revised)

In order to be a justifier, you need to have propositional content, and you need to represent that proposition assertively.

Only Beliefs (Revised)

Only beliefs (or other states that are epistemically like beliefs) represent propositions assertively.

Therefore, only beliefs (or other states that are epistemically like beliefs) can be justifiers. (Pryor, 2005, p. 188)

It is strange that Pryor calls this the Master Argument for Coherentism because, if sound, it would not establish Coherentism. If sound, this argument concludes that only beliefs can be justifiers. But, for a start, this is compatible with forms of Foundationalism! For example, one might hold that *unjustified* beliefs can somehow justify other beliefs. Those other beliefs would be immediately justified since they did not *derive* their justification from the unjustified beliefs; how could they – the unjustified beliefs had none to give. In this case it would still be true that only beliefs can be justifiers, but some unjustified beliefs would play this foundational role. While Pryor seems to recognize this possibility, he nevertheless does nothing to rule it out or to supplement the Master Argument to make it validly imply Coherentism. (Pryor, 2005, p. 188; p. 191)

Another form of Foundationalism compatible with the Master Argument so-presented is one where non-basic beliefs are justified by justified basic beliefs, but the justified basic beliefs are not justified by anything else at all (either their justification is a brute fact, or perhaps they somehow justify themselves). Again, here it would be true that only beliefs can be justifiers, but this theory would maintain that some justifiers could both be justified and justify other beliefs without themselves being justified by anything else.

In order to establish Coherentism, the revised *Content Requirement* and *Only Beliefs* premise need to be supplemented with the following premise:

Justifier: Nothing can be justified without a discrete justifier.

Adding this premise would rule out the two types of Foundationalism I cited above that are compatible with the Master Argument as Pryor presents it. *Justifier* would block either unjustified or justified beliefs from playing a foundational role whereby they justify other beliefs without themselves being justified, or from justifying themselves, since the condition requires that the justified and the justifier be separate and distinct states. Combined with *Content Requirement* and *Only Beliefs*, this argument, if sound, would establish that only beliefs can be justifiers, and since they need to themselves be justified by something else, they would have to be justified by further beliefs.

But this is still not enough to establish Coherentism. *Content Requirement*, *Only Beliefs* and *Justifier* entail that only beliefs can be justifiers, and they too need justifying by other beliefs. But this is compatible with Infinitism. Recall that Infinitism is the view that an infinite regress of justification need not be vicious. An Infinitist alleges that justificatory chains do not merely *transmit* justification, but that if the chain is long enough, it can *generate* justification.³² In order to block this position and to support Coherentism, we need a further premise:

Finite: Infinite justificatory chains cannot generate justification.

We now have a Master Argument for Coherentism. *Content Requirement*, *Only Beliefs*, *Justifier* and *Finite* jointly entail that all justification is mediate, necessarily involves only beliefs, and cannot regress ad infinitum. In addition to these explicit features, the key to characterizing this position as version of Coherentism is to understand

³² Again, see Klein (2005) for details.

the *mediation* relation between beliefs as a matter of coherence or fit. That is, a belief is justified iff it coheres (i.e. ‘fits’) with other beliefs in one’s noetic structure, the noetic structure has a finite number of members, and no member is justified immediately (i.e. no belief derives justification from anything other than a belief, and every belief must derive its justification from somewhere). That is, the justificatory chain includes some closed loops, which results in justification, generated in virtue of the coherence or fit between the belief tokens that constitute the subject’s noetic structure.

If the Master Argument is sound, it establishes how and why immediate justification is impossible – namely because according to theories of immediate justification, some beliefs are justified not by other beliefs, but by *experience*, which this argument concludes is impossible since the only justifiers are beliefs which must be justified by a finite number of other beliefs.

Assessment of the Master-Argument for Coherentism:

A few words about the premises are in order. First, “propositional content” is taken to mean that the content is truth-evaluable. Propositional content is a species of what is sometimes called “representational content”. Such content is characterized as capable of being correct or incorrect, and as being capable of representing veridically. It is a further question whether such content is conceptual. Saying that the content is “assertively” represented is another way of saying it purports to be saying how the world is, unlike other propositional attitudes such as desiring that *p* or entertaining that *p* which make no claim to say how the world is. (Pryor, 2005, p. 187) Second, it is assumed, to

my mind without argument, that beliefs are the only states that represent propositions assertively.

The motivation for the *Content Requirement* is a tendency to think that justifiers need to stand in logical relations to the beliefs they justify. As Davidson famously says in defense of the principle that only a belief can justify another belief, “Sensations [which have no representational content] cause some beliefs and in this sense are the basis or ground of those beliefs. But a causal explanation of a belief does not show how or why the belief is justified.” (Davidson, 2000, p. 157) For example, if a belief is the product of a reliable belief-producing mechanism, that will explain why the subject believes as she does. But on Davidson’s account, unless the subject *believes* her belief is reliably produced, and bases her belief on this belief, the belief will not be justified. Appeal to a cause will explain *why* someone believes as she does, but not if she *ought* to believe as she does. Justification, Davidson holds, is established by appeal to reasons, and the only reasons available to us are our beliefs.

Davidson notes that perception is usually assumed to be epistemically relevant because it is what connects the world to our beliefs about it. (Ibid., p.156) He claims that sensations are thought to be “candidates for justifiers because we often are aware of them. The trouble we have been running into is that the justification seems to depend on *awareness*, which is just another belief.” (emphasis added) (Ibid., p. 157) I deny that conscious awareness is just another belief. Surely there is a distinction to be drawn between being consciously aware that *p* and believing that *p*. The first is merely a kind of awareness (that need not involve belief). One’s experience presents one with assertive representational content – one’s experience purports to say how the world is, and one can

be conscious both of the world and of the nature of one's experience of it. The latter attitude is a belief, viz., an assertive attitude one takes to a proposition. The content of awareness can also include how one's experience *appears* to one; the content of a belief is a *proposition* that purports to represent the way things are in one's perceptual experience. I will argue in chapter II that such non-doxastic awareness is the key to admitting immediate justification that meets epistemic internalist constraints.

I agree with Davidson that things outside our ken are justificationaly irrelevant – but it is a mistake to think our consciousness is limited to beliefs. We are also conscious, or immediately aware, of the world and our experience of it. It is this awareness that gives us justification for our non-inferential perceptual beliefs. Or so I will argue. By denying the *Only Beliefs* premise, and allowing certain kinds of experiences themselves to play a justifying role, we have undermined the Master Argument for Coherentism. Accordingly, we no longer have a reason to doubt that our examples of *prima facie* immediate justification are probative.

The Revised *Only Beliefs* Requirement is False, But Why Accept the Revised *Content Requirement*?:

Pryor notes that many advocates of immediate justification will accept the line of reasoning I gave above and accept experiences as counterexamples to the *Only Beliefs* premise. (Pryor, 2005, p. 189) This move is obviously sufficient to undermine the Master Argument for Coherentism. But what about the *Content Requirement* premise? Pryor says that it is a very demanding premise that is motivated by what he labels the *Premise Principle*:

Premise Principle

The only things that can justify a belief that *p* are other states that assertively represent propositions, and those propositions have to be ones that *could be used as premises* in an argument for *p*. They have to stand in some kind of inferential relation to *p*: they have to imply it or inductively support it or something like that. (Pryor, Ibid.)

The first part of this principle rules out the possibility that beliefs can be justified by non-representational states such as facts about one's reliability, one's sensations and other non-representational mental activity. But that is just what the revised *Content Requirement* demands. It is the latter part of the *Premise Principle* that is more demanding. Specifically, it is more demanding by requiring that the justifiers' propositional content could also be used as premises in an argument. While the *Premise Principle* may motivate the *Content Requirement*, it is not the only motivation, as I will explain below.

The *Content Requirement* and the Egocentric Concept of Justification:

But what about the weaker *Content Requirement*? What motivates it? Pryor considers different possible motivations and finds them wanting. But rather than supporting Coherentism, I submit that the *Content Requirement* is not only true, but it is demanded by an important and valuable conception of justification.

Reflection on our practice of epistemic evaluation reveals that there are different things that we value epistemically. Remember that epistemic justification is minimally a matter of epistemic appropriateness, where the appropriateness of the belief is assessed relative to the truth. There are different ways of making the notion of justification

precise. But there is no reason at the outset to think there is a univocal sharpening that best captures what it is to be epistemically justified.

Analogously, we have the concept of being coloured. Surely there are minimal constraints on what it is for something to be coloured, e.g. it must have the property of producing a sensation in us as a result of reflecting light in a certain way, etc. But there are different ways of being coloured and there is no reason to think that any one colour (red, green, etc.) best captures what it is to be coloured. Similarly, I submit, with epistemic justification, i.e. there are different ways of being epistemically justified and there is no reason to think that any one concept of justification (Deontological, Evaluative, Reliabilist, etc.) best captures what it is to be epistemically justified. If this is correct, it need not be *relativism* about justification (or colour), but rather, a *pluralism*.

However, our way of thinking about epistemic justification ought to be constrained by our epistemic values, interest and goals. Examples of such values include notions of epistemic responsibility and reasonableness that underlie what has been called the Deontological Concept of Justification. A further value is embodied in the Guidance Concept of Justification, which says roughly that an essential feature of justification is that it helps guide the agent in how to proceed intellectually. I do not want to defend these concepts of epistemic justification. However, the kinds of considerations that motivate them are closely related to, and in fact presuppose something that might be called the Egocentric Concept of Justification, conceived of here as concerning:

not the third-person question of whether someone else's beliefs are true or reliably arrived at, but instead the first-person question about the truth (or reliability) of my *own* beliefs [...] [when I] ask whether I ultimately have any good reasons for thinking that my beliefs are true or indeed that they are reliably arrived at. (Bonjour, 2001, p.6; 14)

Given the fundamentally first-person nature of epistemic concepts that relate to such first-person questions, it can be seen that Egocentric epistemic justification must at least in part be a matter of possessing reasons, reasons to which the subject can appeal in wondering if her beliefs are true.

Given this concept of justification, we can begin to see how it motivates the Content Requirement: if we are seeking grounds to think our beliefs are true, the grounds need to be accessible to us. If we want grounds that allow us to defend our beliefs if challenged (even if we raise the challenge for ourselves), the grounds need to be accessible to us. In the next chapter I will go on to explain why if our grounds need to be accessible to us, the Content Requirement must be true. I will argue that, for example, my reason for believing that Jack was at the party is not the photo representing his being at the party, but my *seeing that* the picture represents him as such. My episode of seeing that serves as my reason has propositional content. Similarly one's reason for believing that one has meningitis is not one's aching head, but that fact that one *feels that* one has a severe headache. One's episode of feeling that serves as one's reason also has propositional content.

In the next chapter I will argue that it is a necessary condition of Egocentric justification obtaining that the subject is aware of the grounds for her belief. I will also argue that non-representational states cannot serve as justifiers for the Egocentric concept of justification since we are not able to access them in the right kind of way. After further motivating the necessity of the awareness requirement with what I will call the Subject's Perspective Objection, I will elucidate the nature of this awareness. In so doing I will

respond to the objection that the awareness condition either leads to vicious regress, or else is entirely unmotivated.

CHAPTER TWO

THE AWARENESS REQUIREMENT AND ITS NATURE

2.1. On The Necessity of Awareness:

At the end of the last chapter, I sketched what I am calling the Egocentric Concept of Justification. Given the fundamentally first-personal nature of such an epistemic concept, this kind of epistemic justification is at least in part a matter of possessing reasons to which the subject can appeal in wondering whether her beliefs are true. That these reasons need to be accessible to the subject applies in the first person case where one is assessing one's own justification, as well as when one is evaluating whether someone's else's beliefs are justified in this egocentric sense – their reasons need to be accessible to them as well.³³ Asking the first person questions that characterize the egocentric concept of justification, such as 'what reasons do I have to think my beliefs are true', reveals constraints on not only what kinds of things can count as grounds for belief if it can be justified egocentrically, e.g. certain properties of conscious mental states, but also whether and in what way the subject must have access to the fact that such grounds exist and are indeed adequate for justifying the belief in question. I intend to defend the following as a necessary condition of a belief's being epistemically justified for a subject:

³³ Asking about the justificatory status of beliefs one actually holds may be the central case, but of course one can also ask what one *would* be justified in believing, which can be asked before any belief on the topic is actually held.

Awareness Requirement:

S's is justified in believing that *p* only if

- i) there is something, X, that contributes to the justification of S's believing that *p*; and
- ii) For all X that contributes, S is aware (or potentially aware) of X.

At least three questions about the awareness requirement immediately spring to mind: first, is it really true that awareness is a necessary condition for justification obtaining? Secondly, if awareness is required, what is the nature of this awareness? These will be the primary focus of this chapter. But thirdly, what does 'contribute' mean here mean in the first conjunct? It is not that X *justifies* S's believing that *p*; otherwise, the second conjunct stating the awareness requirement would be superfluous. But it is not superfluous. I maintain that the Awareness Requirement is necessary for justification, and it is met only if both of its conjuncts are satisfied. My focus in this chapter will be motivating and defending ii). However, the force of i) is that not just awareness of *anything* can justify belief. Only certain kinds of things justify certain beliefs. For example, if B is a belief that *p*, viz., the belief that a given contingent empirical proposition is true, then having a perceptual experience as of *p* may justify that belief, whereas hoping or imagining that *p* will not. I will return to this issue and a further defense of i) in chapter 5.

What I do not mean to convey by speaking of a factor 'contributing' to one's justification is that one needs to be aware of only *some* of the justificatory elements, and not others. My view, as it will emerge, is that *all* things on which epistemic justification turns must be such that the subject is aware or potentially aware of them. Accordingly, in this chapter I will first argue that awareness of one's grounds for a belief is a necessary condition for that belief to be epistemically justified. This will be shown through drawing

attention to our considered judgements about cases. I will consider two kinds of cases: first, I will clarify and draw out the implication of cases of clairvoyance and other cases of unusual but reliable cognitive faculties. Secondly, I will show that even many of those who claim to reject the awareness requirement implicitly appeal to it to motivate their own accounts of justification. By considering the arguments for “Mentalism”, I shall show that, unless an awareness requirement is presupposed, the cases adduced by Mentalists are of no intuitive force. In turn, by considering just these cases I will try to establish that the awareness requirement is necessary for justification.

That will then lay the groundwork for giving an account of the nature of this awareness in the latter half of the chapter. By doing so I will examine and respond to an argument against the awareness requirement that we might call Bergmann’s Dilemma. My solution invokes what I call ‘non-doxastic strong awareness’ in a way that avoids the regresses that Bergmann advances.

Before doing so, however, I need to say something about defeaters and the nature of defeat since defeasibility is in play in the cases I rely on to establish that external factors, that is, factors external to the subject’s consciousness awareness, are insufficient for epistemic justification.

Defeaters and the Nature of Justificatory Defeat:

Most basically, a defeater is something that provides one with a reason for ceasing to believe another of one’s beliefs, either by rebutting one’s belief by being inconsistent with one’s beliefs, or by undermining or undercutting one’s reasons for thinking it true.

As I will explain below, depending on the *kind* of defeater in question, it may be a belief, a proposition, or a fact.³⁴

A rebutting defeater is one where the belief is inconsistent with another belief one holds. For example, if one believes that an express train runs on Sunday, but is then told by an official holding an updated train schedule that there is no Sunday service at all, one now has a rebutting defeater for one's belief that an express train will come on Sunday. The second belief that is based on the congruence between the official's testimony and the updated train schedule is inconsistent with one's earlier belief. Since one is unjustified in holding both beliefs at once, one ought to abandon at least one of the beliefs.

Of course one may have a defeater-defeater, a belief that trumps the would-be defeater. Perhaps one thinks that train schedules are notoriously inaccurate and the train employees have in the past proven to be wanton liars. Either way, it is important to note that since a rebutting defeater and the belief it purports to defeat are inconsistent, one would be unjustified in believing both at once.

An undercutting defeater, on the other hand, is not inconsistent with a belief one holds, but it casts doubt on one's reasons for holding the belief. For example, one might believe that if the lights are on inside, the library is open. But if one learns that for security reasons the lights are always left on whether the library is open or not, this new information becomes an undercutting defeater for the belief that the library is open

³⁴ Besides the issue of justificatory *defeat*, there is the related but distinct phenomenon of one's justification being *outweighed* by other reasons one acquires. Suppose one believes that *p* in light of one's evidence *q*. *q* however is not conclusive evidence for *p*, but merely provides one with *prima facie* justification to believe it. Suppose at a later time one learns that *r*. *r* is strong evidence against *p*. But is possible that *r* does not undercut *q* – *q* is still a reason to believe that *p* – nor is *r* inconsistent with *q*, and so it does not rebut it. Instead, *r* *outweighs* *q*. So this is not a case of defeasibility, but it is a case of one's justification being negatively affected by new information that one acquires.

simply because the lights are on. The belief that arises based on the security practices of the library is not inconsistent with the belief that the library is open if the lights are on. However, it does significantly undermine one's certainty of grounds for thinking the library is necessarily open. In light of the undercutting defeater, it would be irrational to have this defeater while adamantly insisting that the library is open solely based on the building being lit.

When considering possible defeaters, it is important to note whether the defeater is intended to rebut or undercut. While rebutting defeaters are intrinsically stronger than undercutting defeaters, in that a rebutting defeater is always also an undercutting defeater (on the assumption that manifest inconsistency is the limit case of 'casting doubt on one's reasons for holding the belief'), a successful undercutting defeater is sufficient to defeat the justification of the belief in question.

Kinds of Defeaters:

Doxastic Defeaters:

There are many different kinds of defeaters. First, one must distinguish between mental state or doxastic defeaters and factual or propositional defeaters. A mental state or doxastic defeater is a proposition that defeats another of one's beliefs by virtue of being *believed*, regardless of whether it is true or not. For example, there is debate whether or not religious pluralism is a defeater for religious exclusivists who maintain that their religious beliefs are immediately justified, without appeal to evidence.³⁵

³⁵ For a defense of the claim that pluralism need not be a defeater for unevidenced Christian belief, see Plantinga (2000), p. 437-457. For a contrasting view, see Robert McKim for an argument that the appropriate response to the facts of pluralism is what he calls the "Critical Stance". This position is partially characterized by a principle to the effect that "disagreement about an issue or area of inquiry

While religious pluralism *may* be a defeater for one who holds the relevant beliefs about the facts of religious diversity, it can hardly be a defeater for those who are unaware of such diversity through no fault of their own (not the result, say, of ignoring the facts) and so who lack the potentially defeating beliefs. For example, isolated tribes or pre-literate people living before the widespread dissemination of information cannot be subject to this kind of defeater if they are unaware of it. In these cases the defeater is only *potential* since the person either lacks the beliefs about the facts of pluralism, or if he does hold relevant (potentially) defeating beliefs, he is unaware of the connection between the defeater and his unevidenced religious belief.

Propositional Defeaters:

A factual or propositional defeater, by contrast, is a proposition that must meet the following conditions: a) it must be true, and b) if it were added to the set of a subject's beliefs, then the resulting set would no longer Warrant the belief in question.³⁶ Key here is that factual defeaters must be true and they need not be believed nor need they be within the easy epistemic grasp of a subject, given the subject's evidence in order to defeat. It is important to note that what is defeated here is *Warrant*, and not justification.³⁷ As I will argue below, this must be the case despite what some of the original authors might have implied by speaking of 'justification'. When, for example, Lehrer and Paxson 1969 and Klein 2000 sometimes speak of these kinds of defeaters

provides reason for whatever beliefs we hold about that issue or area of inquiry to be tentative." (McKim, 2001, p. 141)

³⁶ For early discussions of this kind of defeater, see Lehrer and Paxson (1969) and Klein (2000).

³⁷ Again, Warrant is being understood here in Plantinga's sense of that which converts true belief to knowledge. See 1.1. for more on Warrant and its relation to other epistemic concepts.

defeating ‘justification’, they must be implicitly taking ‘justification’ to be Warrant, for reasons I will give below.

As we will see, factual defeaters defeat *knowledge*, not justification. A standard example of a factual defeater, I take it, is Goldman’s famous fake barn case. (Goldman, 1992a) That example was introduced as a Gettier case to refute the causal theory of knowledge.³⁸ In the barn case one *believes* that one sees a barn, it is *true* that one sees a barn, and one is *justified* that one sees a barn on the basis of one’s perceptual experience – but, one fails to *know* it is a barn since it is in a field full of visually indiscriminable fake barns. The idea is that one does not know one sees a barn because it is just luck that one believed truly as one did; one’s belief could have easily been false on the same evidence (assuming, as some contest, that one’s evidence in this case is one’s perceptual experience as of a barn³⁹). Further, the proposition about the fake barns is true and if it were added to the set of a subject’s beliefs, then the resulting set would no longer Warrant the belief in question. Thus it constitutes a factual defeater.

Note, however, that the intuition is ***not*** that the presence of fake barns defeats *justification*, but rather it is that they defeat *knowledge* (perhaps by defeating *Warrant*, if there is such a thing). First, it is highly intuitive that one has a justified belief about the barn on the basis of one’s perceptual experience of a barn, regardless of what other facts, if any, might obtain, unbeknownst to oneself. Second, if factual defeaters *did* defeat justification, one would no longer have a Gettier case, since while there may be no received definition of a Gettier case, minimally a Gettier case is an instance of a *justified*,

³⁸ The Casual Theory of Knowledge holds roughly that a subject knows that *p* iff the fact that *p* suitably causes the subject’s true belief. For an early expression of this kind of theory, see Goldman (1992c).

³⁹ See ch. 9 of Williamson (2000) for an argument against a phenomenal conception of evidence, and in defense of a factive view that one’s evidence is what one knows. That is, according to Williamson, knowledge (and only knowledge) is evidence.

true, belief that fails to be knowledge. But the fake barn case *is* a paradigm instance of a Gettier case. Therefore, the subject in fake barn country does have justified belief, which he must, since his situation clearly constitutes a Gettier case. Given this, it is not clear how, if at all, merely factual defeat can be extended to justification, at least as we have been understanding the notion.

But can factual defeaters just be defined in terms of *justification*, in the sense that we have been understanding it, rather than being aimed at knowledge? In other words, might we understand a factual or propositional defeater as follows: as a proposition that must meet the following conditions: a) it must be true, and b) if it were added to the set of a subject's beliefs, then the resulting set would no longer *justify* the belief in question? We could of course offer this analysis as a stipulative definition, but the problem is, there does not seem to be a class of defeaters that function in this way.

Take the fake barn case again. Of course there are facts/propositions that if *believed* then the resulting set of the subject's beliefs would no longer justify the belief in question, e.g. the true proposition that the subject is in fake barn country. But this is just an example of a *doxastic* defeater. But as we saw when considering the subject in fake barn country, and whether or not his situation constitutes a Gettier case (which would entail that he is justified in believing as he does), we simply do not judge that facts or propositions that are not believed (and 'should' not be believed, given the subject's evidence) in any way defeat the subject's justification. The fact that we judge that propositional/factual defeaters apply to knowledge, but not justification, offers further support for the kind of epistemic internalism I am advocating, i.e. only that of which the subject is (potentially) aware can contribute or detract from her justificatory status.

Normative Defeaters:

Finally, one can distinguish normative defeaters. These defeat “by virtue of being propositions that S *should believe* given the evidence which is available to S.” (Lackey, 1999, p. 475) Like propositional defeaters, normative defeaters are meant to defeat justification whether they are believed or not. Although unlike propositional defeaters, I hold that normative defeaters do genuinely target justification, rather than just knowledge (or Warrant). However, the existence of normative defeaters that need not be believed to be effective is consistent with what I have said about propositional defeaters above. Because while propositions not believed and outside the easy epistemic grasp of a subject do not defeat justification, normative defeaters *should be believed* whether they are believed or not, but this is so only *given the subject’s evidence*.

In the case of the fake barns above, the potentially (but not actually) defeating facts were outside the subject’s ken, and he could not have been reasonably been expected to discover them, or appreciate their epistemic import, given the evidence he has. This is quite different than the case of Norman the clairvoyant I discuss below, where, given his evidence, in some intuitive but clear sense, he ought not believe as he does. Therefore, at least in the case of Norman, normative defeaters seem to be the kind of defeaters Bonjour has in mind.

Cases of Clairvoyance and Unusual But Reliable Cognitive Faculties:

Laurence Bonjour introduces cases of clairvoyance as arguments against the claim that awareness of one’s grounds is not necessary for justification. (Bonjour, 1985) Bonjour defines clairvoyance as “the alleged psychic power of perceiving or intuiting the

existence and character of distant states of affairs without the aid of any sensory input". (Bonjour, 1985, p. 38) Reliable cases of clairvoyance, if not actual, seem at least to be coherent possibilities. As such, they can help test our intuitive understanding of epistemic concepts.

Bonjour begins by offering three different cases of reliable clairvoyance as putative counter-examples to the thesis that external reliability is sufficient for positive epistemic status. The subjects in these cases, Samantha, Casper and Maud, are all reliable clairvoyants, that is, in the cases in question, it is stipulated that they reliably produce true beliefs via their clairvoyance. The structure of the cases is similar in other ways as well. The subjects have no positive reasons or evidence that they are clairvoyant. Further, in each case the subject is aware of a defeater of some kind.

Take Bonjour's first case:

Case1. Samantha believes herself to have the power of clairvoyance, though she has no reasons for or against this belief. One day she comes to believe for no apparent reason that the President is in New York City. She maintains this belief, appealing to her alleged clairvoyant power, even though she is at the same time aware of a massive amount of apparently cogent evidence, consisting of news reports, press releases, allegedly live television pictures, etc., indicating that the President is at that same time in Washington, D.C. Now the President is in fact in New York City, all evidence to the contrary, being part of a massive official hoax mounted in the face of an assassination threat. Moreover, Samantha does in fact have completely reliable clairvoyant power under the conditions that were then satisfied, and her belief about the President did result from the operation of that power. (Bonjour, 1985, p. 38)

Here the subject has formed a perfectly reliable belief, but she lacks any positive reason for thinking she is capable of forming a reliable belief in the way she did. Further, she is aware of many defeaters that contradict her belief (the news report, etc.) and she has no defeater-defeaters to rebut them. It is because of this that Bonjour concludes that even

though she has a reliably true belief, she does not *know* the President is in New York City because she lacks justification for the belief.

If one sets aside questions of Warrant and knowledge, it seems clear that Samantha lacks *justification* for her belief. This is no doubt in part due to the fact that she is disregarding large amounts of compelling evidence that her belief is false. By believing improperly, Samantha's belief is thereby unjustified. The only complicating factor that might make judging whether or not Samantha is justified is that she does have the background belief that she is clairvoyant, however, it is stipulated that that belief has no support for or against it. But what is clear is that even if that belief alone did somehow contribute to the *prima facie* justification of her belief about the President's whereabouts, she obviously possesses many undefeated doxastic defeaters that are sufficient to defeat or outweigh any justification that belief may have otherwise enjoyed.

The same point can be made using a completely ordinary case. Suppose a highly reliable informant tells me that the President is in NYC. But then fifty other informants tell me that he is in Ottawa. Suppose further that we stipulate that these fifty informants are in fact, unbeknownst to me, utterly unreliable. What am I justified in believing? Surely I have justification to believe that the President is in Ottawa; so if I believe that the President is in NYC, my belief would be quite unjustified. Even if first informant's testimony did contribute to the *prima facie* justification of my belief about the President's whereabouts, I obviously possess many undefeated doxastic defeaters that are sufficient to defeat or outweigh any justification that belief may have otherwise enjoyed.

The next two cases Bonjour considers, those of Casper and Maud, are structurally similar to that of Samantha. In these cases, instead of having a defeater for the specific

clairvoyant belief in question, they have defeaters against the subject himself having clairvoyant powers in the second case, and against *anyone* being clairvoyant in the third case. Here, too, Bonjour concludes that the subjects lack justification because, even though the belief is true and reliably formed, the subject's belief is unjustified since it is held in the face of seemingly cogent defeaters. Here again, if one sets aside the question of whether the subjects have *Warrant* for their belief and hence *know* the proposition in question, it can be agreed the subjects lack *justification* for their defeated belief. Again, even if a dogmatic externalist wants to maintain that the clairvoyants know the whereabouts of the President, it is clear that the clairvoyants are nonetheless unjustified since, as far as they can tell, all available signs point to the falsity of the belief. To be sure, all three cases show that the presence of undefeated-(doxastic and/or normative) defeaters is sufficient to undermine epistemic rationality.

In light of the considerations raised above, some have opted for a qualified externalism, one that contends that reliability is not sufficient for positive epistemic status – it can be defeated. This seems right. Still, a qualified externalism does not require positive evidence that the belief is true or is reliably formed if it is to enjoy positive epistemic status, but it does require that the subject have no good or strong reasons of which she is aware (or should be aware) against the belief in question. This position falls halfway between internalism and an unqualified externalism. The former requires both positive grounds for a belief as well as the absence of undefeated grounds against a belief if it is to have positive epistemic status.

The rationale for requiring positive grounds that motivates the internalist need not require that these grounds be beliefs – I will argue that the demand for positive grounds

can be satisfied by the subject being consciously aware (in a non-doxastic way) of her grounds, which themselves need not be beliefs (for example, in the case of perceptual belief, the grounds may be perceptual experiences of certain kinds). An unqualified externalism says that reliability or some other external factor is sufficient for positive epistemic status, regardless of any evidence (positive or negative) that the subject is aware of. Bonjour's case of Norman is set in terms of a qualified externalism. Bonjour writes,

Case 4. Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable. (Bonjour, 1985, p. 41)

In the case of Norman, he reliably produces clairvoyant beliefs, and in this case there are no doxastic defeaters present. Still, since Norman lacks reasons or evidence for the belief that the President is in New York City, Bonjour concludes Norman's belief cannot enjoy positive epistemic status. Bonjour goes on to reason as follows: if Norman believes he is clairvoyant, that belief clearly lacks positive epistemic status on both the internalist's grounds since he lacks reasons for thinking so, as well as on the externalist's grounds, since the belief that he is clairvoyant is not reliably produced. (Bonjour, p. 42) That is, there is nothing in the case as described that the power of clairvoyance also reliably produces true beliefs about the source of the beliefs. As such, any unevidenced beliefs about the subject's clairvoyance itself would be unreliably produced, and so on the externalist's theory, unjustified. If, on the other hand, Norman does not believe that

he is clairvoyant, from his point of view there is no way he could reliably arrive at the true belief that the President is in New York City. From Norman's own subjective conception of the situation, his clairvoyant belief feels no different than a mere hunch or lucky guess. From this Bonjour concludes that Norman lacks justification and therefore does not know which city the President is in.

Someone might object that Bonjour moves too quickly here by claiming that if the subject is not aware of positive grounds for his belief, he will therefore regard his belief as no better than a hunch. Such an objector would have a point in that an unevidenced clairvoyant belief need not feel like a hunch or lucky guess. But, as I will show below, that is insufficient for the belief to be justified.

Consider three types of cases:

- A) A subject wonders if there is a table in the next room. The subject enters the room, has a conscious perceptual experience of a table, and on the basis of that experience, believes that there is a table in front of him.
- B) A subject wonders if there is a table in the next room. The subject has no idea whether there is a table in the next room, so he guesses. He has no further grounds of which he is aware about whether or not there is a table in the next room. We can stipulate that his guess is unreliable.
- C) A subject wonders if there is a table in the next room. The subject has an inclination to believe that there is, although he has no further grounds of which he is aware about whether or not there is a table in the next room. Here we stipulate that this belief was produced by his highly reliable clairvoyant faculty.

Bonjour seems to treat case B and C the same – if a subject does not have access to his grounds, then from his own perspective, his belief will feel utterly random, as a guess would feel.

The problem with guesses or hunches in this context is that we do not think that they give us any reason to think that a belief that is based on them stands any greater chance of being true. Sometimes we are allowed to rely on hunches and guesses; but not when our concerns are *epistemic*. Sometimes we may have pragmatic reasons that demand that a decision is made quickly where truth is not our (primary) concern, or sometimes we must act without all the relevant information. For example, in a game show one might have pragmatic reasons to guess which door the prize is behind, in which case a guess might suffice. But in this case, the guess cannot serve as the basis for an epistemically respectable judgment about the location of the prize. This is because in epistemic matters, our primary concern is truth. We know that hunches and guesses are not truth-conducive in any sense. They have as much chance of leading to truth as to falsehood and, as such, they cannot contribute to epistemic justification.

What about so-called ‘educated’ guesses? Can they never lead to epistemically justified belief? In some circumstances they can, but they do not do this by virtue of being pure guesses. An educated guess is one that *is* based on grounds of some sort, although the grounds are only partial or incomplete. Similarly, what if a hunch proved to be a reliable guide to truth? In this case, again, this is not a ‘pure’ hunch, but a hunch supplemented with inductive evidence that this is an epistemically appropriate way of forming beliefs in the circumstances. It is clear, however, that without inductive grounds or at least partial evidence, hunches/guesses can never lead to epistemically justified

beliefs due to the fact that they lack any sort of truth-connection. In Chapters 4 and 5 I will further explore what kind of truth-connection must exist between proper grounds and epistemically justified belief.

At any rate, it does not follow from the fact that the subject has a belief produced by his clairvoyance that it needs to strike the subject as being random. It seems there could be cases where the clairvoyant holds his clairvoyant belief with conviction, even strong conviction. In fact, the clairvoyant could take his belief as obviously true, just in the same way that we take our perceptual beliefs formed on the basis of conscious visual experience as obviously true. So, contrary to what Bonjour says, the clairvoyant belief could have the functional properties of normal perceptual belief. It need not present itself as hunches and lucky guesses do. But despite that, the clairvoyant's belief is not justified.

A key difference remains that explains why the unevidenced clairvoyant is not justified, whereas our beliefs based on conscious visual experiences are: in the latter case, we are consciously aware of our grounds. This seems to be the best explanation of these cases. The reason that awareness is justificationaly relevant is the following: whether or not a belief is justified is not a brute fact. Obviously, simply believing something does not necessarily make it true. Even believing something with *conviction* does not make it true. If a belief is justified, it thereby has some connection to truth (the question of the exact nature of the truth-connection is for the moment an open one; again, see Chapter 4 and 5). If a belief is justified, it is justified in virtue of some further thing(s) obtaining. Call this thing that contributes to the belief's justification its *ground*. My position is that

one of the things that it is necessary that one is aware of is one's *grounds*, if one's belief is to be justified.

There is a distinction to be drawn, however, concerning what is meant by 'grounds', as there is a danger of running two important issues together. In the paper "Other Minds", J.L. Austin notes that concerning knowledge and belief, it is common to ask '*how* do you know' and '*why* do you believe', but not '*why* do you know' and '*how* do you believe'. (Austin, 1961, p. 46) Of course one can ask 'why do you know?' if by that one means, roughly, 'why did they tell you?', or something similar. But this is not what is meant by asking 'why do you believe?', which is asking for justification, not an explanation of how the belief was acquired. Related issues arise in a similar way with justified belief. That is, by putting the 'how do you know' question in terms of justification, one might be asking two very different things. On the one hand, one might be asking what the *source* of one's justified belief is, that is, in what way did one come to acquire this belief? On the other hand, one might be asking what the justificatory basis of one's belief is, that is, what reasons does one have for thinking it true?

These two questions obviously can come apart. For example, suppose one is justified in believing that bachelors are unmarried men. What are one's grounds for this belief? As to the sources question, perhaps one's way of coming to believe this is that one was taught the relevant concepts as a child, which seems to be a paradigm example of an empirical belief source. This is to assume that the five senses are empirical sources of belief. One's justificatory basis for one's belief about bachelors, however, is non-empirical. Presumably, one's justification arises from conceptual understanding, or some other *a priori* justifier, e.g. in virtue of the self-evident character of the proposition.

It is possible for one to be aware of the source of one's belief but not its justificatory basis, as the case above shows. Similarly, it is possible for one to be aware of one's justificatory basis for a belief while not being aware of its source (perhaps because one acquired the belief long ago and has since forgotten how it was acquired). The issue of memory presents an interesting special case for the awareness requirement that I will not be able to address in this thesis, but will leave for further research. However, it is *acquiring* justified belief that is the central case of epistemic justification, not its mere retention. I place special emphasis on the acquisition of belief, rather than its retention, because it is conceptually more fundamental: one cannot retain a belief unless one has acquired it; but one can acquire a belief but fail to retain it (perhaps because one has poor memory). In the paradigm case of acquiring justified perceptual belief, I contend that one must be aware of *both* the source of one's belief, as well as its justificatory basis, which together constitute the ground of one's belief. Without awareness of one's ground, from one's own perspective the justification of one's belief is a brute fact, which we know is inconsistent with its being justified (more on this below).

For example, in the case of a belief about the external world based on a conscious perceptual experience, the subject is aware of her grounds, so she is aware of a reason to think her belief is true. From her perspective her belief 's being justified is not a brute fact: it is justified by her experience. She is aware of both the source of her belief (sensory perception) as well as its justification (that is, the fact that the conscious perceptual experience has a certain phenomenal character and an intentional content). It is the clairvoyant's lack of awareness of his grounds that makes his belief unjustified. The clairvoyant is neither aware of the justificatory basis of his belief nor of its source,

and either of these lacks of awareness is sufficient for the belief to be unjustified. As to the sources question, the clairvoyant lacks awareness of any way that he could come to have justified beliefs about the whereabouts of the President. He also lacks awareness of any reason to think that what he believes is true (independently of the question of how he acquired this reason).

There is a further difference between the way conscious awareness contributes to justification in the normal visual perception case, and the way the lack of awareness ensures that the clairvoyance based belief is unjustified. In the normal case the subject is sensitive to kinds of defeaters that the clairvoyant is not. In the case of belief about the external world based on a conscious perceptual experience, the subject is aware whether the belief is consistent with her grounds. If she loses the source of her belief, or if the belief's justification is undermined or rebutted by further considerations, the belief will be dropped. That is, in an ideally rational person, a belief is "extinguished" when that person judges that belief not to be supported by appropriate reasons. (Cf. Scanlon, 2000, p. 20) A clairvoyant belief is not sensitive to these kinds of defeaters. The clairvoyant belief cannot be dropped in response to considerations about its source or justification since the clairvoyant is not aware of the belief's grounds. This insensitivity to grounds implies that the belief is not justified, and further, that such sensitivity is made possible only by an awareness of grounds.

These considerations give rise to what we might call the Subject's Perspective Objection, which along with the egocentric concept of justification, is the primary motivation for the awareness requirement respecting justification. Michael Bergmann introduces the Subject's Perspective Objection like this:

Bergmann's Subject's Perspective Objection: If the subject holding a belief isn't aware of what that belief has going for it, then she isn't aware of how its status is any different from a stray hunch or an arbitrary conviction. From that we may conclude that from her own perspective it is an accident that her belief is true. And that implies that it isn't a justified belief. (Bergmann, 2006, p. 12)

Given what I have argued above, it is clear that this formulation needs revising. As I argued, the subject may not regard a belief whose grounds she is ignorant of as a mere hunch or arbitrary conviction. My revised formulation below takes this into account. We can now state the Subject's Perspective Objection thus:

Revised Subject's Perspective Objection (SPO): If the subject holding a belief is not aware of the belief's grounds, then she is not aware of how its status is any different from a stray hunch/arbitrary conviction/strong but groundless conviction. As such, if the subject regards her belief as justified, she ought to regard the fact that the belief is justified as a brute fact. From that we may conclude that from her own perspective she ought to regard it as an accident that her belief is true, i.e. from her perspective a belief whose justification appears to be a brute fact could just as easily be false. And that implies that the belief is not justified.

While qualified externalism is surely an improvement over the unqualified varieties, it does not go far enough. If there is a rationale for needing the absence of defeaters for a belief to be rational, it is the same reasoning that motivates the need for positive evidence as well – the difference between these two types of evidence is one of degree, not kind. If a belief is unjustified if it is held in the face of evidence against it, even if unbeknownst to the subject the belief is true and reliably formed, the belief is still unjustified if it is held and the subject has absolutely no indication that it is true.

A subject ought not to just think that her belief is justified in virtue of something or other, when she is not aware of it, or what it could be. She will have to say that she

has no reason to regard her belief as justified as we saw with some of the clairvoyants above, and so she cannot then just maintain, in an *ad hoc* way, that she is justified in virtue of something that she is not currently aware of. This is at best a mere promissory note, with no assurance of repayment. But perhaps she has a *reason* or *evidence* that in certain circumstances if she takes herself to be justified, then she usually is justified. But then in such a case the awareness condition is met: what justifies her in belief is not something that she is not aware of. Rather, her belief is justified in virtue of her justified belief that if she is in certain circumstances and believes something on topic Φ , then she is usually justified, *and* that she is in those circumstances.

For example, Jack believes that the Battle of Hastings was in 1066. Here are two possible scenarios: he may not recall where he learned this fact, and if he merely thinks that he is justified in virtue of something or other, but he is not aware of what it is, then his belief is unjustified, as the SPO rightly predicts. Take the second scenario: Jack might be aware of inductive evidence, that combined with a justified belief about what kind of proposition he is considering, might yield justification in believing that proposition. That is, Jack may have strong inductive evidence that he is usually correct about dates of events in early medieval English history, and he may be aware that he is now considering a proposition about an event in early medieval English history. As such, his belief may well be justified, and if it is, it is justified by internalist lights.⁴⁰

⁴⁰ Cases of putative justified belief that might be thought to be problematic for the awareness requirement, since it not obvious what the subject's grounds might be, include 'bare' memory, the *a priori*, knowledge of our own mental states, the positions of our own limbs, etc. While these are interesting cases that a complete theory of epistemic justification should ultimately address, I will not be able to address these cases in this thesis. Instead they must be left for further research, although this passage indicates one possibility for how to address certain cases of memory justification in a way consistent with epistemic internalism. Nevertheless, my primary focus will be on the justification of perceptual belief. For more on the scope of this thesis, and my rationale for this focus, see the Introduction.

But to be sure, from the subject's point of view, beliefs that emerge from a qualified externalism appear to have no better epistemic status than at best a strong but groundless conviction. Therefore, it seems that the common sense case to be made for understanding epistemic justification surely favors the internalist, conceived here as someone who accepts the awareness requirement as necessary for justification.

Before moving on to discuss Mentalism, it is important to note a feature of the dialectic. So far I have argued that certain cases are not cases of epistemic justification. This is an intuitive verdict, which I hope is found plausible. In addition to not being instances of epistemic justification, all of the cases I considered above lack a certain awareness condition. It is of course a further claim that this is *why* they lack justification. I submitted that this is the best explanation of our judgments about the cases, and accordingly, I concluded that the awareness condition is necessary for epistemic justification obtaining. But given that I am arguing by inference to the best explanation, there will always be a gap in the argumentative strategy since such arguments employ nothing like deductive validity. Despite this intrinsic feature of such arguments, I hope I have done enough to bridge the gap, and thereby have shown that the awareness requirement is necessary for epistemic justification.

2.2. Mentalism:

So far I have argued that the awareness requirement is a necessary condition for justification by considering cases of clairvoyance (the same kinds of judgments would be evoked by considering cases of unusual but reliable cognitive faculties generally, respecting the grounds of which the subject has no awareness). I will now show that even many of those who claim to deny the awareness requirement implicitly appeal to it to motivate their accounts of justification. By considering the arguments for “Mentalism” I will show that, unless an awareness requirement is presupposed, the cases that such arguments appeal to are of no intuitive force.

Mentalism is a theory of epistemic justification that is least like the other positions that have been called internalism that I have discussed so far. In fact, if to be an internalist one must endorse *some* kind of awareness requirement like the one I have been suggesting, then Mentalism would be regarded as thoroughly externalist. Richard Feldman and Earl Conee, the chief advocates of Mentalism, claim that Mentalism is the thesis that “a person’s beliefs are justified only by things that are internal to the person’s mental life.” (Feldman and Conee, 2001, p. 233) So the “internal” in internalism here means internal to the mind.

Specifically, their official formulation of Mentalism is as follows: “S: The justificatory status of a person’s doxastic attitudes strongly supervenes on the person’s occurrent and dispositional mental states, events, and conditions.” (Ibid., p. 234) From this formulation, they express the main implication of S thus: “M: If any two possible individuals are exactly alike mentally, then they are alike justificationally, e.g., the same beliefs are justified for them to the same extent.” (Ibid.) This latter claim simply spells

out a consequence of supervenience. But it is worth noting that the antecedent may be impossible to fulfil. That is, it is not clear if two individuals can be *exactly* alike in respect of their mental states. For example, if one subject believes that he is *f*, and the other believes that he (the second subject) is *f*, the subjects seem to be in different mental states, since they believe different things. Accordingly, the supervenience formulation may not illuminate anything. I will further address the issue of internalism/externalism about mental content and its possible implications for epistemic internalism in chapter 3.

Feldman and Conee expressly say that an advantage of Mentalism is that it parallels what philosophers have called “internalism” in the philosophy of mind and ethics. (Ibid., p. 235) In the philosophy of mind, the internalist claims roughly that the contents of one’s thoughts supervene upon one’s intrinsic physical properties – the outside environment and relations to it have no bearing on the content of one’s thoughts. But of course unless Mentalists are also internalists about the mental, they need not be committed to the position that epistemic justification depends on internal psychological properties, since these positions are distinct. Besides being both a Mentalist and an internalist about the mental, one could also be a Mentalist about justification and an externalist about the mental; or an internalist about the mental and an externalist about epistemic justification; or of course, an externalist about both. The position I am defending is an internalism about justification, not in the Mentalist’s sense of course, but rather conceived as someone who accepts the Awareness Requirement. For the moment, I am remaining agnostic about whether externalism or internalism about mental content is correct (more on this issue in chapter 3).

In ethics, an internalist is someone who holds that one has a reason for action only if one has some desire that would be satisfied by performing that action. The position is internalist since desires, which are internal states of affairs, are the only thing relevant to what reasons for action one has. Nevertheless, it is unclear why calling a position in epistemology “internalism” is an *advantage* simply because the label is used in similar ways in other areas of philosophy. What label we use is irrelevant to the truth of the position.

It is important to note that Mentalism as Feldman and Conee construe it has no access or awareness requirement of any kind. However, there may be *some* connection between Mentalism and access. If one endorses an access requirement, depending on what it is that one holds must be accessible (e.g. only certain properties of mental states are relevant for epistemic justification), then a kind of internalism that is characterized by an awareness requirement may entail Mentalism. But of course it need not, since one might hold that the things one must be aware of are not themselves mental. For example, in direct realist accounts of perception, subjects are aware of things in the world – they are not mental.

Be that as it may, the converse is not true: Mentalism does not entail forms of access internalism since Mentalism only requires that the factors that determine justification are internal to the mind, not that the subject can tell that they are internal, or even detect their presence, wherever they happen to be located. Whether the “mental states, events and conditions” lie below the surface of consciousness and are thus undetectable is justificationaly irrelevant for the Mentalist, given their official formulation of their position, something the access theorist flatly denies.

Arguments For Mentalism:

In defending Mentalism, Feldman and Conee construct six pairs of cases where they invite the intuition that in the first instance the subject has a justified belief and in the second case the belief is intuitively not justified, or else one belief is more justified than the other. Feldman and Conee contend that the best explanation of these apparent epistemic differences is that there are “internal” differences in their preferred sense of internal to the mind, although not necessarily accessible or within the subject’s “epistemic perspective”, as I maintain is necessary for justification. If this is an intuitively correct diagnosis of the cases, this is meant to support their thesis that differences in justification have an entirely mental origin, and have nothing to do with other factors, including what the subject is aware of. Their first two cases and comments on them are representative of all the cases they provide and their assessment of them.

They introduce them as follows:

Example 1 Bob and Ray are sitting in an air-conditioned hotel lobby reading yesterday’s newspaper. Each has read that it will be very warm today and, on that basis, each believes that it is very warm today. Then Bob goes outside and feels the heat. They continue to believe that it is very warm today. But at this point Bob’s belief is better justified.

Comment: Bob’s justification for the belief was enhanced by his experience of feeling the heat, and thus undergoing a mental change which so to speak “internalized” the actual temperature. Ray had just the forecast to rely on.

Example 2 After going out and feeling very warm, Bob goes back in and tells Ray of the feeling. Here are two versions of the relevant details:

2a) Bob is in fact a pillar of integrity, but Ray has no reason to think so. As far as Ray can tell, it is just as likely that Bob is trying to deceive him as that Bob is telling the truth.

2b) Bob is a pillar of integrity, and Ray has observed and recalls many examples of Bob’s honesty and none of dishonesty.

In example (2b) Ray's belief that it is very warm becomes more strongly justified after he hears from Bob. In example (2a) hearing from Bob does not affect the strength of Ray's justification for his belief.

Comment: Bob's honesty, something out of Ray's ken in (2a), has become "internalized" by Ray in (2b). Bob's integrity made no justificatory difference to Ray's belief until it was suitably brought into Ray's mind. (Feldman and Conee, 2001, p. 236)

By considering these cases, it is far from clear whether they effectively distinguish Mentalism from a theory of justification that embraces the awareness requirement, since the feeling of heat in example 1, and observations and recollections of Bob's honesty in example 2 are both things that the subject is aware of. If we stipulate that he is NOT aware of them, the cases lose their intuitive plausibility. While Feldman and Conee *may* be right about which subjects are justified and to what extent, I deny that they offer the best explanation of those facts.

For example, *if* in example 1 Bob's belief is better justified when he goes outside, this is because he is *aware* that he experiences the heat. The mere experience of heat, without the subject's awareness of it (if such a case is possible), would not yield justification on its own. Incidentally, it is somewhat ambiguous in what way Feldman and Conee maintain that Bob's justification is *enhanced* over Ray's upon going outside and feeling the heat. Surely he has an additional source of justification, that is, he now has his experience as well as the testimony of the newspaper. So in one sense his justification is more stable since it now has multiple independent sources. His belief is justificationaly "overdetermined", so to speak. But that is not necessarily to say that the *strength* of his justification is increased. At any rate, if we stipulate that Bob has the

experience of the heat, but is not aware of it, it is far from clear that his justification differs in *any* way from Ray.

Similarly, in the second set of examples, the emphasis is on what Ray can *tell* about Bob's honesty. It is Ray's *consciously recalling* inductive evidence that Bob is honest that marks the difference between 2a) and 2b). If instead Ray had those same inductive grounds stored in *unconscious* memories, it is not at all clear that there are justificatory differences between case 2a) and 2b).

Feldman and Conee are seeking to give the best explanation of our verdicts about these cases. They claim the explanation is that there is "no justificatory difference to Ray's belief until it [the ground for the belief] was suitably brought into Ray's mind." (Feldman and Conee, 2001, p. 236) That is surely right. What Feldman and Conee fail to recognize, however, is that being "suitably" brought into mind means that the subject is *aware* of his grounds. If it is possible for grounds to be in the mind but for the subject to be unaware of them, it is false that the subject is justified, as careful consideration of these cases reveal. This shows that even those who explicitly deny the awareness requirement in the course of their defense of their theory of justification implicitly rely on it. Feldman and Conee cannot show that the Mentalist condition is properly fulfilled without the awareness requirement because all the cases they discuss possess the awareness requirement. If we explicitly exclude the awareness requirement, the cases lose any plausibility they had.

Therefore, while the Mentalist is right that there can be no justificatory difference without a mental difference, there is crucially an additional necessary condition that needs to be met, namely the Awareness Requirement. So besides supporting Mentalism,

the arguments offered by Feldman and Conee offer support my view as well. Granting that Mentalists obviously do not think that just *any* mental difference yields a justificatory difference, they still fail to put any constraints on what kind of mental differences are relevant to epistemic justification. My account of justification specifically addresses this problem by arguing for the necessity of the Awareness Requirement. Accordingly, my theory of justification better explains Feldman and Conee's examples than Mentalism does.

So far I have argued that awareness of one's grounds for a belief is a necessary condition for that belief being epistemically justified. Without it, an account of justification gives rise to the Subject's Perspective Objection. This was shown through drawing attention to our considered judgments about cases. I considered two kinds of cases: cases of clairvoyance and cases that were originally offered in defense of Mentalism. By considering the latter cases, I showed that even those who claim to deny the awareness requirement must implicitly appeal to it to motivate their accounts of justification. Having made a case that the Awareness Requirement is necessary for justification, I will now turn to give an account of the nature of this awareness.

2.3. On the Nature of Awareness: Responding To A Dilemma for Internalism

In a recent book, Michael Bergmann argues that the Awareness Requirement that defines the kind of epistemic internalism I favour generates a dilemma: either it leads to a vicious regress, and thereby to a scepticism about whether any of our beliefs are justified, or it is entirely unmotivated (i.e. that meeting such an awareness condition cannot meet the SPO, the primary motivation for the awareness condition). As a result, Bergmann argues, epistemic internalism, the Awareness Requirement, and the intuitions that motivate it captured in the SPO must be abandoned, despite their initial plausibility. By critically examining Bergmann's dilemma we can see why it fails, and in so doing clarify and defend what kind of awareness that epistemic justification demands.

Bergmann's Dilemma:

Bergmann presents his argument against the awareness requirement (and thereby internalism) as follows:

A Dilemma for Internalism

- I. An essential feature of internalism is that it makes a subject's actual or potential *awareness* of some justification-contributor a necessary condition for the justification of any belief held by that subject.
- II. The awareness required by internalism is either *strong* awareness or *weak* awareness.
- III. If the awareness required by internalism is strong awareness, then internalism has vicious regress problems leading to radical scepticism...
- IV. If the awareness required by internalism is a weak awareness, then internalism is vulnerable to the SPO, in which case internalism loses its main motivation for imposing the awareness requirement.
- V. If internalism either leads to radical scepticism or loses its main motivation for imposing the awareness requirement (i.e. avoiding the SPO), then we should not endorse internalism.

VI. Therefore, we should not endorse internalism. (Bergmann, 2006, p. 13-14)

‘Strong’ and ‘weak’ awareness are introduced as technical terms. Strong awareness involves “*conceiving* of the justification-contributor that is the object of awareness as being in some way relevant to the justification or truth of the belief”. (Bergmann, 2006, p. 13) Weak awareness, by contrast, involves no such conceiving. (Ibid.) I accept premise I since it merely expresses my Awareness Requirement, and I will accept premises II and V for the sake of argument. Also, I will argue that Bergmann is correct about premise IV for the reasons he gives; accordingly, I will respond to the dilemma by seizing its other horn – I will show that Bergmann is wrong that strong awareness leads to vicious regress (premise III). By doing so I will vindicate the Awareness Requirement and the constraint on epistemic justification it imposes.

Being ‘weakly’ aware of one’s ground for a given belief involves either i) not conceiving of the ground in any way, i.e. a non-conceptual awareness of the ground; or ii) while the ground is conceptualized *as* being a certain way, it is not conceived *as* relevant to the truth or justification of the belief it is meant to support. An example of non-conceptual weak awareness is the one experienced by, say, a cat watching a game of chess. The cat, presumably lacking the concept of chess, is seeing a situation in which chess is being played, but it does not see *that* a game of chess is being played. The cat is aware of chess being played by being consciously related to a state of affairs, but it is only weakly aware.

Conceptually weak awareness, on the other hand, consists of the subject being aware of her ground, conceptualizing it as being a certain way, but failing to be aware of its relevance to what is believed. For example, say Jill believes that Jack is in the room.

She also sees someone in the room who is in fact Jack dressed in a Santa Claus costume handing out presents to children. She is consciously aware of the person appearing a certain way. Further, she conceptualizes it as such. However, while she is aware of Jack's presence in the room (because she sees him in a Santa outfit), she is not aware *that* Jack is in the room, since he is cleverly disguised in the costume. Jill fails to draw any connection between her believing that Jack is in the room and her conscious experience of what is in fact Jack in the costume. She fails to conceive of her perceptual experience *as* relevant to the justification of what she believes. Upon reflection, Jill is not aware of any reason to think that her belief that Jack is in the room is true.

Consider one final case to illustrate the difference between this kind of awareness and the strong awareness introduced above. Suppose someone simply looks in a room and seems to see that John is there and believes that John is there. Is this an instance of conceptually weak awareness? A reason to think so is that the subject need not be thinking (in any relevant sense) in terms of evidence or justification. The reason that the case as described so far is a case of *weak* awareness is that it is possible for the subject to have both the experience and the belief and not *base* the latter on the former, and in no way even *potentially* appreciate an evidential connection between them.

What needs to be the case in order for this example to be an instance of strong awareness is the addition of a further state: the subject is aware that the justification contributor (here, his visual experience of John's being in the room) is relevant to the truth or justification of what is believed (namely, that John is in the room). What this reveals is that strong awareness takes the form of a second order state. What one is aware

of, or at least potentially aware of, is an evidential 'fit' between one's grounds on the one hand, and what is believed on that basis, on the other. More on this below.

As Bergmann points out, if the subject is merely weakly aware of the grounds of his belief, e.g. non-conceptually aware or conceptually aware of the grounds, but not aware of any connection between the grounds and a supported belief, then that subject is vulnerable to the SPO. From the subject's own perspective, it is an accident that what he believes is true, since from his own perspective, the status of his belief is no better than a hunch or arbitrary conviction, or as I grant, at best a strong but groundless conviction (which is nevertheless incompatible with a belief being justified). But since the SPO is given against externalism, if it applies here as well against weak awareness, this would undermine the Awareness Requirement.

Bergmann makes two further distinctions regarding strong awareness. First, strong awareness may be doxastic or non-doxastic. Doxastic strong awareness "involves the *belief* that the object of awareness is in some way relevant to the truth or justification of the relevant belief." (Bergmann, 2006, p. 14) Non-doxastic strong awareness is simply strong awareness that does not involve belief. But since it must involve the concept of justification, how can it involve such concepts without belief? It is true that such states of awareness are conceptually laden, i.e. the subject must have certain concepts, such as justification and truth to be in the state, but the state is not itself a belief. Compare this to a conceptually laden perceptual experience: a subject must have certain concepts, such as a 'cat' and 'mat' for example, in order to be in the state of seeing that the cat is on the mat -- but here the state itself is not a belief. Rather, it is a

kind of experience. Similarly for non-doxastic strong awareness. I will develop this idea further below.

Secondly, strong awareness may be *actual* or *potential*. As the label suggests, strong actual awareness requires that the subject actually be aware, “whereas the potential strong awareness requirement demands only that the subject be able on reflection alone to be aware.” (Ibid.) If the awareness was non-propositional in form (a kind of weak awareness), then there could be actual awareness that only potentially gives rise to a state of justified belief. Such cases would be neither quite actual nor potential cases in the above sense. But as we saw in the Jack and Jill case above, non-propositional or other kinds of weak awareness are not justificationaly relevant. What is required is specifically strong awareness, which may in turn be either actual or potential. Bergmann argues that all versions of strong awareness lead to objectionable regresses, and hence scepticism.

Take the obvious regress first generated by this principle:

Actual Doxastic Strong Awareness Requirement (ADSAR):
S’s belief B is justified only if (i) there is something, X, that contributes to the justification of B and (ii) S is actually aware of X in such a way that S justifiedly believes that X is in some way relevant to the appropriateness of holding B. (Bergmann, 2006, p. 14-15)

By demanding a further justified belief (as stated in (ii)) as a necessary condition of having any justified beliefs at all, it is easy to see how both a *mental state* and a *complexity* regress are generated. (Bergmann, 2006, p.16) As Bergmann points out, for S’s belief B to be justified, it requires a further justified belief with respect to a further justification contributor (call it X1). This belief will be P1: “P1: X1 is in some way relevant to the appropriateness of holding B.” (Ibid.)

But this justified belief will, of course, require a further justified belief about its grounds (call that ground X2). This belief will be P2: “P2: X2 is in some way relevant to the appropriateness of believing that X1 is in some way relevant to the appropriateness of holding B.” (Ibid). The result here is that if ADSAR were true, we would require an infinite hierarchy of beliefs of ever-increasing complexity.

The fundamental objection here is that doxastic strong awareness generates a *complexity* regress, not a *mental state* regress. This can be seen if one makes the move from requiring actual to merely potential access. Here, an advocate of the position could argue that there is merely a *potential* infinite regress, not an actual infinite regress, so the regress is harmless, assuming that we could have an infinite number of dispositions to believe. (Bergmann, 2006, p. 16) While this move may block the mental state regress, it still requires us to *be able* to believe propositions of infinite complexity. Since most of us cannot grasp even the third and fourth level propositions, given our cognitive limits, it is obvious that we are unable to grasp the infinitely complex. Accordingly, this theory predicts that we do not have any justified beliefs. Given that this is intuitively false, strong doxastic awareness (actual or potential) must be rejected. Or so Bergmann claims.

A Reply:

Consider again the problematic formulation of the strong access requirement that occurs in the requirement’s second conjunct:

Actual Doxastic Strong Awareness Requirement (ADSAR):
S’s belief B is justified only if (i) there is something, X, that contributes to the justification of B and (ii) S is actually aware of X in such a way that S *justifiedly believes* that X is in some way relevant to the appropriateness of holding B. (emphasis added) (Bergmann, 2006, p. 14-15)

What generates the regress here is that it is assumed that the kind of strong awareness required is *doxastic*. The solution to Bergmann's dilemma is to realize that there are forms of non-doxastic strong awareness that can contribute to a subject's having non-inferential or immediate justification. To defend this claim, it is important to understand both what immediate justification and non-doxastic strong awareness are. In 1.5. I defended the claim that there is immediate justification, that is, justification that does not come from holding other justified beliefs. There I showed that putative examples of immediate justification *prima facie* support the existence of immediate justification, and undermined the cogency of a central argument against them (i.e. the so-called Master Argument for Coherentism). With that in place, we can turn to non-doxastic strong awareness.

On Non-Doxastic Strong Awareness:

With an idea of what is meant by immediate justification, we can return to the idea that one's non-doxastic strong awareness allows for immediate justification in a way that avoids the regresses that Bergmann advances. Recall that the motivation for the Content Requirement (see 1.5.) is a tendency to think that justifiers need to stand in evidential relations to the beliefs they justify. As we saw in chapter 1, Davidson says in defense of the principle that only a belief can justify another belief, "Sensations [which have no representational content] cause some beliefs and in this sense are the basis or ground of those beliefs. But a causal explanation of a belief does not show how or why the belief is justified." (Davidson, 2000, p. 157) For example, if a belief is the product of a reliable belief-producing mechanism, that will explain why the subject believes as she

does. But according to Davidson, unless the subject *believes* her belief is reliably produced, and bases her belief on this belief, the belief will not be justified. Appeal to a cause will explain *why* someone believes as he does, but not if he *ought* to believe as he does. Justification, Davidson holds, is established by appeal to reasons, and the only reasons available to us are our beliefs.

Recall further from chapter 1, that Davidson notes that perception is usually assumed to be epistemically relevant because it is what connects the world to our beliefs about it. (Ibid., p.156) He claims that sensations are thought to be “candidates for justifiers because we often are aware of them. The trouble we have been encountering is that the justification seems to depend on *awareness*, which is just another belief.” (emphasis added) (Ibid., p. 157) Bergmann seems to think of awareness in a similar way as consisting largely only of beliefs, or if he does think that there are forms of awareness that are distinct from belief, e.g. he allows for a certain kind of non-conceptual ‘weak awareness’, he seems to think that they all share a property in common with belief that is problematic in this context, viz., the property of being apt for epistemic justification (which threatens to give rise to vicious regress).

Crucially, I deny that conscious awareness need just be another belief. Surely there are distinctions to be drawn between experiencing that *p*, believing that *p*, and the further state of being aware that one is in these states, and the relations between them. The first is obviously a kind of *experience* that can be shown to be distinct from belief. For example, suppose one has a conscious perceptual experience of the Muller-Lyer illusion. In this case, one has an experience as of the lines appearing to be of unequal length. Having such an experience does not entail that one *believes* that they are unequal

in length. If one is familiar with the illusion one will believe that the lines *are* the same length, despite their appearance. This state is a belief, viz., an assertive *attitude* one takes to a proposition.

In addition to these two states, one can be aware that one is in them, and certain facts about them. The content of this awareness is how one's experience *appears* to one, both in the sense of an awareness of the state's representational content (i.e. one's experience purports to say how the world is), as well as *how* or in what *way* the content is presented, e.g. as a perceptual experience with a certain phenomenal character, as opposed to a hoping or a wishing. The content of a belief is just a *proposition* that purports to represent the way the world is, and one can be aware that one is in such a state while simultaneously being aware that one's experience is such, and is relevant to what is believed.

Just as it can be shown that perceptual experience and perceptual belief are distinct, it can be shown that the kind of awareness I am interested in is distinct from those states as well. Obviously this kind of awareness is distinct from *perceptual experience* since some lower creatures have perceptual experiences without this kind of awareness. First, the kind of awareness I am pointing to is a second order kind of ability that they lack, since they have only first order mental abilities.

Secondly, this kind of awareness is *conceptual* in that one cannot be in it unless one possesses the constituent concepts required to represent the content of what one is aware. A kind of perceptual experience that I was considering above is also conceptual in this sense. For example, just as the cat cannot see *that* a game of chess is being played without the concept of chess, the cat cannot be aware *that* he is having a visual

experience of that kind which supports a belief of his without possessing various fairly sophisticated concepts, e.g. chess, the concept of a perceptual experience, belief, etc. We of course possess such concepts, and as such this problem does not arise for us. But since lower order animals lack these kinds of concepts, they do not have the capacity for this kind of awareness. Since they lack this awareness, but do have a capacity for a kind of perceptual experience, the two are distinct.

This kind of awareness is also distinct from *belief* since someone can have a perceptual belief and a perceptual experience while lacking awareness of the relation between them, perhaps because they are distracted or otherwise inattentive. This is a kind of ‘weak’ awareness that is too weak to play a justificatory role, as I alluded to above in the formulation of Bergmann’s Dilemma. On the other hand, a subject might have *awareness* without belief, since she may have a defeater for her belief and rationally drop it, while still being aware of her perceptual experience, its *prima facie* evidential import to a belief she held, etc. On the face of it, there is no reason to assume that this state of awareness that remains is itself a belief.

The fact that this kind of awareness is a kind of experience, rather than a belief, can be appreciated by noting that awareness, like other kinds of experiences, cannot be defeated. That is, suppose one seems to see a wombat in the garden and also believes that there is a wombat in the garden. If a doxastic defeater is introduced, e.g. the subject justifiably believes that he has been slipped a wombat image-inducing hallucinogen, and the subject is rational, he will suspend judgement on whether or not there is a wombat in the garden. What gets defeated is the subject’s belief, viz., he no longer believes that there is a wombat in the garden. But his perceptual experience remains unchanged. New

beliefs do not affect the way things look to him. Similarly, the subject's awareness that he is having a certain perceptual experience, and its relation to what was believed, is not, and it seems cannot, be affected by defeaters – unlike belief. Strong awareness, therefore, is distinct from belief, since the former cannot be affected by defeaters, whereas the latter can.

These considerations seem to show there is a kind of conscious awareness that is distinct from perceptual experience and belief. So my contention is that one can be simultaneously aware that: one's experience is such, and, that one believes that things are as one's experience represents them as being, as well as that the grounds stand in an evidential relation to what is believed. Such awareness is an instance of 'strong' non-doxastic awareness, and it (or its possibility upon reflection) is necessary for such a belief to be justified if the SPO is to be avoided. Of course one *may also* believe these things (i.e. *believe* that one is having a perceptual experience, *believe* that one's grounds support what one believes, etc), in which case the question of how these higher order beliefs are justified arises, and a regress is generated. But again, there is no reason to assume that awareness *must* take a doxastic form.

But if one is strongly aware that these things are so, then does one not *take them to be so*? In which case, in what sense is the state non-doxastic? But this is to assume that the only way of 'taking things to be so' is in the form of belief. But obviously this does not follow. Things are 'taken to be so' when, roughly, they are represented assertively, which is to say, when they are represented as being a certain way. And belief is certainly one such state that functions in this way. But as we already noted in chapter 1, perceptual experiences also represent things as being a certain way, and they are

certainly non-doxastic. Similarly, I submit, with strong non-doxastic awareness. Thus, the following general inference is invalid:

- a) A state is conceptually laden and represents its content assertively
- b) Therefore, this state is necessarily a *belief*.

But perceptual experience is a counterexample to the above principle. There is no general reason, therefore, to assume that strong awareness *must* take a doxastic form, simply in virtue of its being conceptual and representing its content assertively. And no specific reason against the possibility of non-doxastic awareness seems forthcoming either.

Besides conceptually distinguishing strong awareness from one's belief and one's conscious experience, and noting that there is no reason to assume that such a state must take a doxastic form, can more be done to show that it actually exists? I hope to have motivated the need for such states because without it, one is left with 'weak' awareness, which I argued is too weak to play a justificatory role, since it gives rise to SPO. But without further positive argument, invoking strong awareness as a way of blocking Bergmann's dilemma may seem like introducing a *deus ex machina*. Accordingly, I will now offer two arguments for the existence of states of strong awareness, which I will call the Argument From Introspection and the Transcendental Argument, respectively.

First, and most directly, one can detect one's states of strong awareness through introspection. When one attends to one's visual experience of the kookaburra in tree, and one's belief that that is a kookaburra, one can also be aware that one is in these states, as well as the evidential relation between them, i.e. that this experience is good grounds for this belief, since the latter can be justified on the basis of the former. If through

introspection one detects the state, and my earlier discussion functions as an elimination argument, i.e. the state is not a belief, it is not the first order experience, etc., then the state can be taken at face value: a state of strong awareness; a conscious experience that one is having a certain perceptual experience that justifies a belief that is based on it.

But is this empirically implausible? Would one normally say that one wanders around the world aware of one's own experiences and beliefs in this way? Admittedly one does not often attend to these experiences, and as such they tend to go unnoticed. But that is not to say that they are not there and that we cannot turn our attention to them if needed, e.g. in the course of defending the justification of one of our perceptual beliefs, if only to ourselves. It is an unreasonable expectation that if our mental life exhibits certain features, then these need to be constantly consciously manifest to us. Instances of strong awareness are no different in this regard.

For example, we have countless beliefs, the majority of which are not occurrent at one time. Of course we can call many of them to consciousness. But that is not to form the belief, but merely to recall it. The fact that the belief was not occurrent until recalled does not show that it was not held. Even among mental states that one is actually conscious of at a time have aspects that go unnoticed. For instance, many of these states have phenomenal properties, even though we do not attend to these aspects of the states. Consider again when one is conscious that the kookaburra is in the tree, but one fails to take notice of particular shade of its plumage until one attends to it. Of course we can identify such properties, and attend to them through introspection. But it is not as if these properties are *created* through introspection; they are just noticed through introspection, even though in a sense one was conscious of them the entire time.

Similarly, states of strong awareness exist and can be detected through introspection, and this fact is not diminished by the fact that these states are not ones that we often attend to in the course of the daily lives. This in turn does not detract from the fact that they are necessary for epistemic justification to obtain and allow for a kind of rational defensibility if challenged.

Besides the above argument from introspection, I will also offer a transcendental argument for the existence of states of strong awareness. My general methodological approach throughout this thesis has itself been broadly transcendental: I have been arguing that given that we *do* have justified beliefs, what is a necessary condition of our beliefs enjoying such a property? To that end I have been defending an awareness requirement. The following argument is similarly transcendental in that I am enquiring into the necessary conditions of experience. To fix terms, what I mean here by a transcendental argument is one of the following form: since *p* is a necessary condition for *q*, and since *q* is beyond dispute, therefore *p* is the case.

What I take to be obvious and beyond dispute in this context is that we are justified in believing many, if not most, of the things that we take ourselves to be justified in believing (see Introduction). What I also hope to have shown so far is that external factors (such as the reliability of a process) are *insufficient* for epistemic justification to obtain.⁴¹ A consequence of this is that the most radical forms of epistemic externalism are false. To that end I have argued that a form of internalism is true, namely one that embraces an awareness condition. That *some* form of conscious awareness is necessary for epistemic justification to obtain is an intuitive verdict, and one that I hope is also not

⁴¹ See chapters 3 and 4 for arguments that external factors (such as the reliability of a process) are not *necessary* for epistemic justification to obtain, either.

in dispute in this context. What remains at issue is what the nature of this conscious awareness consists of. To that end, consider the following argument:

Transcendental Argument for Strong Awareness:

- i) Some kind of awareness is a necessary condition for epistemic justification obtaining.
(see section 1 of this chapter)
- ii) Such awareness is either 'weak' or 'strong.' (see above)
- iii) Weak awareness is too weak to be relevant for justification, since the SPO would still apply.
- iv) If the SPO applies, then one does not have epistemic justification.
- v) We do have epistemically justified beliefs.
- vi) Therefore the SPO does not apply. (from iv and v)
- vii) Therefore the kind of awareness that must be present is not 'weak'. (from i, iii, iv, v and vi)
- viii) Therefore the kind of awareness that must be present is 'strong'. (from ii, vii)

Since we do have justified beliefs, and awareness is required, but it cannot be weak awareness, then strong awareness *must* exist. That is, strong awareness is necessary for the possibility of our having justified beliefs, and given that we obviously do have justified beliefs, strong awareness must exist. If it did not, scepticism would follow. But scepticism is false. Hence a transcendental argument for the existence of strong awareness.

It is worth noting that in mounting my transcendental argument I am arguing for a more modest conclusion than traditional transcendental arguments tend to aspire to. Traditional transcendental arguments tend to try to refute scepticisms of various kinds by

moving from some fact about our mental life that is beyond dispute to the truth of some extra mental proposition, e.g. Kant famously argues in the *Critique of Pure Reason* that, among other things, certain truths about space and time are necessary for the possibility of experience. I am not attempting to bridge the gap between the mental and extra mental by arguing that facts about the latter are necessary to account for obvious facts about the former.⁴² More modestly, I am merely arguing that it is a necessary condition of being in one mental state (having a justified belief) that we must be in another (having strong awareness); the latter is indispensable to the former.

The kind of strong awareness I am pointing to, like perceptual experience and beliefs, has a propositional content, in that its content is capable of being true or false. Again, my contention is that one can be simultaneously aware that: one's experience is such, and, that one believes that things are as one's experience represents them as being, as well as that the grounds stand in an evidential relation to what is believed. The kind of awareness I am highlighting, on the face of it, seems to be a second order kind of experience. Metaphorically, it is like "seeing that" something is the case. When one is having an experience as of hands, and one believes that one has hands, one can be aware that (or "see that") one's belief is evidentially supported by one's experience. It is difficult to say what *kind* of experience this awareness is since one can be aware that so many radically different *kinds* of things are so, e.g. perceptual experiences, that one's own mental states are such, etc. This may suggest that the awareness involved is a kind

⁴² To bridge *that* gap on the basis of a transcendental argument seems difficult at best, and impossible at worst. As Barry Stroud has argued, it seems that at best a transcendental argument could establish how things must *appear* or what we must *believe*, rather than how they must *be*. He goes on to argue that this gap can only be closed by accepting idealism or verificationism, neither of which seem like very plausible positions in their own right, let alone as responses to scepticism. See Stroud (2000). By not moving from the mental to the extra mental, my argument does not face this difficulty, since there is no gap between merely appearing to be strongly aware and being strongly aware. If I *seem* to be aware that I am having a certain perceptual experience and that it justifies a belief of mine that is based on it, then I am so aware.

of *sui generis* mental state that shares many properties in common with other kinds of experience, without being identical to any one of them.

A concern needs to be addressed. Someone with Bergmann's sympathies might object that if awareness has propositional content, its epistemic status is an open question. However, certain kinds of *perceptual* experience are propositional, in that the nature of such experiences is that it appears *that p*, and that is presumably enough to allow it to stand in logical and evidential relations to beliefs. But we do not think perceptual experience would need, or even admit of, epistemic justification. Similarly, why think that *awareness that p* would need, or could even admit of, epistemic justification? Intuitively, epistemic justification is a property of *beliefs*, not of propositional attitudes or experiences generally.

It seems odd to talk of my *fearing that p* or *hoping that p* or my *experiencing that p* as being epistemically justified. For example, it seems a category error to say things like 'My smelling that the toast is burnt is epistemically justified' or 'I have epistemic justification for my hearing that the car pulled into the garage'. The best explanation of this oddness is that smelling and hearing, as kinds of experiences, are not apt for epistemic justification. Similarly for *awareness*, given that it is a kind of experience, as I argued above. Of course, *beliefs* about one's fear or hope or experience admit of epistemic justification, and those other propositional attitudes and experiences themselves **may** admit of other types of justification, such as moral, pragmatic or aesthetic. But epistemic justification, properly so-called, applies only to beliefs.

I agree with Davidson that things outside our ken are justificationaly irrelevant, but it is a mistake to think our consciousness is limited to beliefs. We are also conscious,

or strongly aware, of our experience. It is these experiences that give us grounds for the justification of our non-inferential perceptual beliefs, combined with our non-doxastic awareness of them. Remember that, given the fundamentally first-person nature of the epistemic concept we are interested in, what interests us are reasons to which the subject can appeal in wondering if her beliefs are true, and thereby goes a step towards avoiding the SPO: if we are seeking grounds to think our beliefs are true, the grounds need to be accessible to us. If we want grounds that allow us to be aware from our own perspective of the non-accidental nature of the truth of our beliefs, the grounds need to be accessible to us in a certain kind of way. Non-representational states cannot serve as justifiers for the egocentric concept of justification, since while one might be able to aware *of* these states (a kind of weak awareness), one cannot be aware *that* these obtain in way necessary to avoid the SPO, as I hope to have shown above.

But now does one need to be aware of the awareness of the connection between sense experience and belief in order for *that* awareness to not fall prey to the SPO? That is, if a second level awareness is needed, is a third level of awareness needed, and if so, is a vicious regress not thereby generated? If not, why not, given the need for the extra level of awareness in the first place? To answer this possible objection, it is important to recall exactly what the SPO says:

Revised Subject's Perspective Objection (SPO): If the subject holding a belief is not aware of the belief's grounds, then she is not aware of how its status is any different from a stray hunch/arbitrary conviction/strong but groundless conviction. As such, if the subject regards her belief as justified, she ought to regard the fact that the belief is justified as a brute fact. From that we may conclude that from her own perspective she ought to regard it as an accident that her belief is true, i.e. from her perspective a belief whose justification appears to be a brute fact could just as easily be false. And that implies that the belief is not justified.

The SPO motivates the idea that when one is enquiring about the *justificatory* status of *belief*, then a certain kind of awareness is required if that belief is indeed justified. As I argued above, ‘weak’ awareness is too weak for this purpose, and so a form of ‘strong’ awareness is needed. But to demand awareness of this state of awareness on the grounds of the SPO is to commit a category error. The SPO arises at the level of belief, and is relevant to whether or not that belief is justified for the subject. A state of strong awareness, by contrast, is neither a belief (it is a kind of *experience*), nor is it even apt for epistemic justification (again, given that it is a kind of *experience*), and so the SPO is inapplicable.

I hold that to be justified in believing that *p* on the basis of a perceptual experience with the content *p*, one must be (or potentially be) aware that the experience is justifying evidence for one’s belief. I argued that such a state is necessary for *epistemic justification* to obtain. But it could hardly be a necessary condition of having an experience of a certain kind in general that one would need a further experience of it. If so, one could never have any experiences at all, since presumably having an infinite number of experiences is impossible.⁴³ Just as having a first order perceptual experience does not require a further state in order for it to be one of which one is consciously aware, similarly the kind of state that constitutes one’s strong awareness of one’s ground and its relation to what is believed can be had without requiring further awareness of it. In short,

⁴³ Perhaps allowing that a subject could have an infinite number of dispositions to have the requisite experiences could block a mental state regress? Perhaps. However, the fundamental problem here is a *complexity* regress. While a subject can be aware of certain conscious states that she is in, most of us do not have the capacity to be aware of many levels higher than that, since the content is too complex to grasp. Cf. p. 101-102 above for a discussion of mental state v. complexity regresses cast in terms of *belief* rather than awareness.

while the kind of non-doxastic strong awareness I advocate is motivated by the SPO when considering whether a perceptual belief is justified on the basis of a given perceptual experience, it is unmotivated, and surely unnecessary to require a third level awareness (or more) either as a condition of having strong awareness, or for strong awareness to play its essential justificatory role.

Someone might object that there is a difference between a belief having the *property* of justification, as opposed to *showing* or *proving* that one's belief has justification. (Pryor, 2005, p. 194) The objection would hold that to conflate having justification and showing justification would simply be a version of what is called a "level confusion", e.g. when one confuses being justified in believing that *p* versus being justified in believing that one is justified in believing that *p*.⁴⁴ I am aware of this distinction and I am not making this confusion. When we reject Norman's belief about the President's whereabouts, for example, we are rejecting it *for the reason* that he has no access to grounds for his belief. We are not questioning his reasons to *think* that he is justified, we are questioning his justification *simpliciter*.⁴⁵

A further constraint imposed by the Awareness Requirement is that one can only base beliefs on what is consciously accessible to one, i.e. beliefs or experiences that are conscious, or were conscious and can be easily brought to consciousness again. I am suggesting that it is a necessary condition of basing a belief on a ground that one appreciates the evidential relevance of the ground to the belief. The result is that

⁴⁴ See Alston (1989d) for a discussion of level confusions in epistemology.

⁴⁵ Richard Fumerton makes a similar point in defense of his Principle of Inferential Internalism. He defines the Principle of Inferential Internalism as follows: "To be justified in believing one proposition *p* on the basis of another proposition *E*, one must be 1) justified in believing *E* and 2) justified in believing that *E* makes probable *p*." (Fumerton, 1995, p.36)

believing for a reason is “discriminative”. As Audi notes, in the case of one belief being based on another, if one’s supported belief is discriminative,

the belief that *r* does not *qua* reason belief, or *basis* belief, as we may call it, tend to give rise to or sustain just any belief, but only those that *S* takes *r* to support; secondly, where *r* is the *only* reason for which *S* believes *p*, so far as *S* is disposed to explain or justify his belief that *p* (e.g. on being asked why he believes it), he i) spontaneously tends to appeal to *r*, and ii) does not spontaneously tend to appeal to other beliefs, in the explanatory or justificatory attempt. (Audi, 1993a, p. 242)

Likewise, if one’s basic perceptual belief is based on a conscious perceptual experience, that experience will justify some beliefs and not others, namely beliefs one takes the experience to support (where this ‘taking’ is strong non-doxastic awareness). Where the experience is the only reason for the belief, one will tend to spontaneously appeal to one’s experience as the basis of one’s belief if asked to justify one’s belief. Such a relation ensures that our justified beliefs are not mere effects of their grounds, but rather are *guided* by our grounds.

On the model of internalist non-inferential justification that I am proposing, my justification for my belief that *p*, where *p* is a perceptual belief about the external world, is a perceptual experience with the content *p*. It is an essential feature of that experience, if it is to count as a justifier, that in addition to it I have (or potentially have) strong non-doxastic awareness of my grounds. My conscious awareness of the experience is directly and immediately available to me simply in virtue of having the perceptual experience. As a perceptual experience, it can stand in justificatory relations to beliefs without it being apt for epistemic justification, since it is not a belief. Similarly, my non-doxastic strong awareness of the perceptual experience satisfies the internalist’s condition, without

being apt for epistemic justification either, since it is not a belief, and therefore no justificatory regress (complexity or mental state) is generated.

But as I introduced with the Awareness Requirement, there is more to being epistemically justified than mere awareness: there must actually be something that contributes to one's justification that one is aware of. Only certain kinds of things that one is aware of will justify certain kinds of beliefs (e.g. imagining that I will win the lottery does not justify my belief that I will win the lottery). In the case of perceptual belief, that ground is obviously a perceptual experience of a certain kind. What properties this ground must have if it is to function as a justifier is a question to be addressed throughout the rest of this thesis. What I hope to have shown here, however, is how internalists can avoid a fundamental dilemma that threatens vicious regress. Our beliefs are justified since we are (or can be upon reflection) strongly aware of reasons to think our beliefs are true in a way that blocks Bergmann's Dilemma since this awareness is *non-doxastic*. In other words, my belief is egocentrically justified since I am non-doxastically strongly aware, from the first person perspective, of a reason to think my belief is true in a way that avoids the SPO and blocks Bergmann's dilemma.

CHAPTER THREE

THE CASE OF THE NEW EVIL DEMON: IS (ACTUAL) RELIABILITY NECESSARY FOR JUSTIFICATION?

3.1. Introduction:

In Chapter 2 I introduced Laurence Bonjour's cases of clairvoyance as a way of showing that external factors are insufficient for justification. Similar style arguments have been offered to show that external factors, such as the reliability of the mechanism supporting the belief, are not necessary for justification. Hence, justification is an entirely internal matter.

The main argument in support of the conclusion involves cases of comparing what constitutes justified belief in the actual world with one's counterpart who lives in a demon world, like the one entertained in Descartes' First Meditation.⁴⁶ As the New Evil Demon problem was first introduced, it was claimed that the demon world is one which by hypothesis is, from our perspective, just like the actual world. Different versions of the thought experiment might be stronger or weaker, that is, the demon may deceive the subject about different ranges of beliefs. For example, in the most radical versions the demon hypothesis, the demon might threaten not only the justification of our perceptual beliefs, but also our justification for seemingly self-evident truths, like Descartes' Cogito, simple logical or analytic truths, or beliefs about the current contents of our own minds.⁴⁷

⁴⁶ I take it that the 'old' and 'new' Evil Demon problems involve the same thought experiment, but draw different conclusions from it. The 'old' problem is the sceptical one, viz., how do we know that we are not dreaming or are the victims of an evil demon that is radically deceiving us? The 'new' problem concerns justification, viz., do our counterpart demon victims differ from us in epistemic justification?

⁴⁷ Jonathan Schaffer argues that there is a demon with unlimited sceptical range who threatens universal doubt. Rather than deceiving us with falsities, the debasing demon, as he calls it, would allow us true

A more modest evil demon scenario is one where what we experience and believe is just like in the actual world, except the demon ensures that all our perceptual beliefs are false. It is this kind of demon scenario that will be the primary focus of this chapter. If sound, the argument based on comparing counterpart subjects in this demon world and the actual world purports to show that external factors are not necessary for justification. Further, and crucial for our purposes here, the argument also purports to show that the two subjects are justificationaly equivalent, and so internalism about justification is supported, as I will explain.

As I will discuss below, many philosophers who have been moved by arguments presented by externalists about mental content deny that such a case is possible. That is, they deny that there could be a world where a counterpart has all the same beliefs that we do, but that all of their beliefs are false (because given externalism about content, what beliefs the counterparts have may be different). Since many people hold that some form of content externalism is obviously true, if it is incompatible with epistemic internalism, this would seriously threaten my internalist position. However, I hold that the necessity of the awareness requirement that I introduced in chapter 2 remains substantially unaffected no matter which view of content turns out to be correct. What is key to the thought experiment is that the two worlds are completely subjectively indistinguishable from each other by those who inhabit them in all the ways that they are consciously aware of, which is what matters for epistemic justification. I will argue that even if a standard moral of the New Evil Demon intuition is untenable due to considerations arising from content externalism, the case can be reinterpreted to support the awareness

beliefs, but meddles with the basing relation in a way that undermines justification. For details see Schaffer (manuscript).

requirement on justification that I introduced in chapter 2 in a way that is wholly compatible with content externalism.

In the final part of the chapter I consider other reasons offered by John McDowell to reject the New Evil Demon intuition. I consider McDowell's position since he is not a paradigm epistemic externalist (in fact, according to Neta and Pritchard (2007), he is a kind of epistemic *internalist*) and yet he rejects the New Evil intuition. If his case is convincing, it will both undermine a motivation for the kind of internalism I endorse, plus it will undermine a central claim of mine, namely, that external factors are not necessary for epistemic justification obtaining. I will try to show that McDowell's arguments leave my position unscathed. That is, factors external to the subject's conscious awareness, such as the *actual* reliability of the belief forming mechanism used in a world, are not necessary for epistemic justification to obtain.⁴⁸

⁴⁸ Even if *actual* reliability is not required, might *conditional* reliability be necessary for epistemic justification to obtain? For example, even if the subject's beliefs are not reliably formed in the actual world, perhaps due to the meddlings of an evil demon, must they need to be reliable in some 'normal' world if they are to be epistemically justified? I will explore and reject proposals of these kinds in chapter 4.

3.2. What is the “New Evil Demon” Intuition?

Stewart Cohen, who originally introduced the New Evil Demon problem, notes that if Reliabilism about justification is true, (i.e. a belief is justified iff it is the product of a reliable belief-forming process), then in the case of the demon world, perceptual beliefs will never be justified since they are always false. (Cohen, 1984, p. 281) Of course a Reliabilist can maintain that some false beliefs can be justified, since while false, they may still be the product of a process that is *generally* reliable. Even so, as a New Evil Demon case is set up, the relevant class of the subject’s beliefs are *always* false, as the demon ensures, and so they are unreliably produced, and hence unjustified by Reliabilist lights.⁴⁹⁵⁰

Cohen argues as follows:

Thus we can imagine two inhabitants of this [demon] world, A, who is a good reasoner, i.e., reasons in accordance with the canons of inductive inference, and B, who engages in confused reasoning, wishful thinking, reliance on emotional attachments, guesswork, etc. Since the beliefs of A and B are both produced by unreliable processes (the evil demon sees to this), a reliabilist theory of justification must render identical epistemic appraisals of both sets of beliefs. Plainly, this cannot be correct. A’s beliefs are conditioned by the evidence whereas B’s beliefs are not. A is a good reasoner whereas B is not. A’s beliefs are

⁴⁹ But even if such processes are not reliable in the actual world in which they are deployed, are they not reliable in *some* sense? Might they be *conditionally* reliable in the sense that they are reliable in some ‘normal’ world, and hence resulting beliefs are epistemically justified? Again, I will explore and reject proposals of these kinds in chapter 4.

⁵⁰ Reliabilism and the issue of false justified beliefs give rise to another potential complication. Suppose that a subject validly infers some beliefs from others. Can these beliefs not be justified, even for the reliabilist, since certain patterns of inference are reliable? Of course in these cases they may generate falsehoods since the premises are false, but the patterns employed themselves are in general reliable. What consequence does this have for epistemic justification?

While some beliefs arrived at in this way **may** be justified, it is hard to say in general what is true of such cases, since there does not seem to be a straightforward connection between valid inference and justification. Valid inference is neither necessary nor sufficient for justification. It is not necessary: many of our justified beliefs are justified non-inferentially, e.g. beliefs in certain conceptual truths, beliefs about our own bodily sensations, etc. It is not sufficient: simply because one believes that *p*, and one believes that if *p* then *q*, it does not follow that one is justified in believing that *q*. Sometimes the rational thing to do is to believe that $\sim p$ instead. As Gilbert Harman has argued, rules of logical implication are one thing, and rules of reasoning are another. See Harman (1986).

reasonable whereas B's belief are not. There is a fundamental epistemic difference between the beliefs of A and the beliefs of B. But the Reliabilist does not have the theoretical means to display this difference. I would claim that the distinction between the beliefs of A and B is marked precisely by the concept of justified belief. (Cohen, 1984, p. 283)

Cohen anticipates what the Reliabilist could say about justification here: while the beliefs of A in the demon world may be “rational” or “reasonable”, this does not affect her justificatory status. (Ibid.) The subject in the demon world has no justification because, on the Reliabilist's account, she lacks reliably true beliefs, which they take as a necessary condition for justification. So even if A and B in the demon world did have the odd true belief, for example, it would fail to be justified, let alone constitute knowledge, because from an epistemic point of view, the truth of the belief was accidental (a kind of demon-created Gettier-case). Cohen rejects this Reliabilist move, saying that “reasonable” and “rational” are “virtual synonyms” for “justified”, so that if one concedes that the subjects are rational and reasonable, one is also thereby conceding they are justified too. Cohen is right that intuitively the subject A in the demon world does have justified beliefs, and we should hold on to this conclusion unless we have a good reason to give it up. What remains at issue is if externalists of different stripes (e.g. about mental content, or justification) can give us such a reason.

A further moral is often drawn from cases like the above in defense of epistemic internalism. Just as subject A and B are not equally justified, even though they are equally unreliable, it is often claimed that subject A, the good reasoner, is just as justified as his good-reasoner counterpart in the actual world. Call A's good-reasoner counterpart in actual world A*. Intuitively, A and A* are equally justified in believing as they do. If so, again actual reliability and other external factors are not necessary for justification

since the two subjects are equally justified and the external factors are present in one case and not another.

I agree with what we can call the New Evil Genius View (NEG), which Ram Neta and Duncan Pritchard define as follows: “(NEG) The extent to which S is justified at t in believing that *p* is just the same as the extent to which S’s recently envatted duplicate is justified at t in believing that *p*.” (Neta and Pritchard, 2007, p. 381) As they rightly point out, this is a deeply rooted intuition that counts in favour of epistemic internalism. They add the qualifier “recently envatted” as way of trying to skirt the problematic issues of content externalism.⁵¹ Depending on the kind of content externalism endorsed, this qualifier may be sufficient to handle the problems I will discuss below. If not, as I will explain, the intuitive support of the New Evil Demon case is not lost in support of epistemic internalism as I conceive of it (by embracing the awareness condition) even if the most austere forms of content externalism are true.

⁵¹ But does their qualifier do that? It depends on the details of the content externalism endorsed. However, demonstratives might be thought to pose a special problem here nevertheless. Suppose that a subject S believes that *that* object is red. His twin S* cannot have that same thought since he is hallucinating. So there is at least one proposition that S is justified in believing that S* is not. Pritchard and Neta’s formulation does not avoid this potential problem. But as I will argue below, such results are compatible with the kind of epistemic internalism I endorse. Key here will be how “sameness” of justification between twins is understood.

3.3. Content Externalism:

Before it is possible to see if epistemic internalism and content externalism are incompatible, it is necessary to understand what these theories amount to. In chapter 2 we saw that epistemic internalism should be understood as the view that holds that the awareness condition is necessary for justification, viz.,

The Awareness Requirement: S's is justified in believing that *p* only if

- i) there is something, X, that contributes to the justification of B; and
- ii) For all X that contributes, S is aware (or potentially aware) of X.

Recall that in Chapter 2 I motivated and defended ii), as well as considered the nature of such awareness. I will address i) more fully in Chapter 5. What I did stress in Chapter 2 was that by a factor 'contributing' to one's justification I do not mean that one needs to be aware of only *some* of the justificatory elements, and not others. My view, as it continues to emerge, is that *all* things on which epistemic justification turns must be such that the subject is aware or potentially aware of them. And what kinds of things can a subject be aware of that are justificatorially relevant? As I will argue in this chapter, only those things which a subject would share in common with his envatted duplicate, as the New Evil Demon case reveals.

Content internalism/externalism are theories about the nature of intentional mental content (what one's thoughts are about), and how it is fixed or determined. How the views are best formulated turns on other views one might hold in the philosophy of mind, e.g. the status of Physicalism (roughly, the view that everything, and therefore the

mind, is in some sense ‘physical’).⁵² In formulating the views as generally as possible, one would want to include traditional content internalists who are not Physicalists. But for simplicity and ease of exposition, I will formulate the views in terms that presuppose the truth of Physicalism.⁵³ Of course the definitions can be altered to make them more inclusive if desired.

So assuming the truth of Physicalism, we may put some of the possible views as follows:

Content Internalism-P: The content of an agent’s mental states supervene upon his intrinsic physical properties.

Is content externalism the denial of internalism? If so, it is far from obvious. While holding that content does not supervene on intrinsic physical states, someone may consistently deny that content is determined by external factors. Perhaps content is some *sui generis* feature alongside the subject’s intrinsic physical states and the external world. Accordingly, we should understand the denial of Content Internalism-P as follows:

Non-Internalism-P: At least some of the content of an agent’s mental states *fails* to supervene upon his intrinsic physical properties.

Besides merely denying content internalism, Content Externalism makes a positive claim that includes some determination relation, such as supervenience, to the subject’s environment. Accordingly, we may understand it thus:

Content Externalism-P: At least some of the content of an agent’s mental states fails to supervene upon his intrinsic physical properties, but rather, is essentially dependent on his physical environment (including his social environment).

⁵² The formulation of Physicalism is not a trivial matter. For scepticism that the doctrine can be formulated in way that is both true and informative, see Crane and Mellor (1990). For a more optimistic line on saying precisely what Physicalism amounts to, see Jackson (1998) ch 1.

⁵³ I will indicate this presence of this presupposition by affixing ‘-P’ to the labels of the positions.

Fortunately for my purposes here, I do not have to take a stand on the truth of Physicalism, or what the exact relation is (if any) between a subject's thoughts and the external environment. The view that Physicalists, non-Physicalists, Content Externalists and Non-Internalists can all agree upon, which threatens to be incompatible with the kind of epistemic internalism I endorse, can be put as follows:

Non-Internalism (General): At least some of the content of an agent's mental states fails to supervene upon the phenomenal features of his mental states.

That is, what a subject's thought is about (viz., what proposition he is entertaining) does not depend upon how things seem to him. According to this view, which most of my opponents will accept, two subjects can seem to be aware of all the same things, and yet the contents of their thoughts differ. If this is true, this might seem to threaten my interpretation of the New Evil Demon case and undermine its support for the kind of epistemic internalism that I am defending.

Content externalism is often presented as obvious on the basis of a style of thought experiment first introduced by Hilary Putnam, and later developed by Tyler Burge. (Putnam, 1973; Burge, 1979) In Putnam's examples, we are asked to imagine two worlds that are *nearly* identical in every respect. The only difference is that on our planet, the stuff in lakes and rivers that we call "water" is H₂O, whereas on Twin earth, the stuff the inhabitants call "water" actually has the molecular structure XYZ – otherwise the two substances are indistinguishable. We are asked to suppose further that the year is 1750 so that molecular structures have yet to be discovered.

Now Oscar, an inhabitant of earth, has beliefs and other attitudes about water, which he expresses in language with the word "water". His twin, Twin Oscar, Putnam suggests, has beliefs and other attitudes NOT about water, since water is H₂O, but rather

about XYZ, (which we can call twater), even though he expresses these thoughts in his language with the word “water”. So the claim is that when Oscar says, for example, ‘I want a drink of water’, he is not saying the same things as Twin Oscar when he says ‘I want a drink of water’. Support for this claim comes from noting that since when Oscar utters the above remark he says something true, and when Twin Oscar utters it he would be saying something false if by “water” he meant water (since the word “water”, stands not for H₂O, but XYZ in his world). Therefore, they cannot be asserting the same proposition. Putnam’s own conclusion about his case is that the “meaning” of terms, which he takes to be their “extension” (roughly, a term’s extension is the set of things it applies to) is not determined purely psychologically, but depends on how the world is. Hence Putnam’s famous dictum, “meaning’s just ain’t in the head”.

The kind of externalism here is generally taken to not be limited to just meanings, or a purely linguistic externalism, but to thought contents. The Twin Earth thought experiments are taken to support content externalism generally, understood here as the view that two different thoughts can be had by internally identical subjects. So despite being internally identical, Oscar has a water (H₂O) thought; Twin Oscar has a twater (XYZ) thought, which, while very similar, are not the same thought, since their truth-conditions differ. They have different beliefs about the stuff found in lake and rivers since *what* it is they believe, i.e., the content of their belief, is the content it is because of the environments the subjects find themselves in.

Tyler Burge presents two arguments for this stronger conclusion as the best interpretation of Putnam’s cases in his paper “Other Bodies”. (Burge, 1982) The conclusion of the first argument is that if Twin Oscar could genuinely have a water

thought, rather than a twater thought, it would be genuinely mysterious how this is possible since he has never had any contact with water, or even with anyone who has. As Burge points out, no one on Twin Earth even uses a word that means water, since their word means XYZ. (Burge, 1982, p. 109) The driving premise of this argument is that for someone to think that p , where p is a proposition, it is a necessary condition of such a thought that the subject possesses the concepts that constitute the proposition. For example, one cannot believe that the cat is on the mat unless one has the concepts of 'cat' and 'mat'. If concepts are the 'building blocks' of thoughts, then without concepts, one lacks the building materials to construct thought. It is difficult to see how empirical concepts like water could be acquired in non-empirical ways. Surely no one would insist that Twin Oscar could have a belief about water because he possessed the concept of water innately!

Still, what I remain unclear about is whether empirical concepts being acquired empirically is supposed to be a necessary truth, or whether it is mere contingently true about beings who are constituted as we are. For example, is it possible that by some freak occurrence (such as being struck by lightening, or being exposed to a kind of radiation) Twin Oscar could acquire the water concept even though there is none in his environment? If so it would be an "empty" concept, in the sense of having a meaning but having no extension? Or it would have neither a meaning nor an extension?

I am raising this question to see what the commitments of externalism are. While empirical concept acquisition on the face of it may seem to pose special problems for externalists, these problems may be no different for internalists about content. That is, acquiring concepts need not in general happen by perceptual encounter with samples.

Many of us have the concept of a mongoose, yet few of us have encountered one. Also, externalism itself allows that someone could have a concept innately, or acquire it from God, so long as the subject was situated in a suitable environment. I will have to leave these issues unresolved here.

Burge's second argument that thought contents are individuated externally is that if they were not, then a large number of Twin Oscar's beliefs would be false. That is, if Oscar has the water concept, then any of his beliefs about the stuff in his world that fills the lakes and rivers would be false since he would believe falsely of XYZ that it was water (which is necessarily H₂O). Burge argues that since "their beliefs were acquired and relate to their environments in exactly parallel and equally successful ways", it is implausible to hold that many of Twin Oscar's beliefs are false." (Ibid. p. 10) I am not sure whether this argument actually *establishes* content externalism or merely *presupposes* it, but it does seem a nevertheless counter-intuitive result, if correct, that if Twin Oscar had the concept water, his beliefs about twater would be false.

The conclusion of this second argument can be strengthened if one realizes that a Twin Earth philosopher could reason in ways parallel to Burge and Putnam. He could argue that if the subjects on Earth (which he would call 'Twin Earth') had the twater concept, then many of their beliefs about H₂O would be false. They would use the words 'I believe that stuff in the ocean is water', but they would be believing falsely since their word 'water' refers to the concept XYZ and there is no XYZ in their world. All they are acquainted with is H₂O. If one is drawn to the conclusion that Twin Earthians have mostly true beliefs about the liquid in their lakes and rivers and we on earth have mostly true beliefs about the liquid in our lakes and rivers, then since the truth-conditions of the

beliefs of the respective world inhabitants are different, this might be thought to offer some evidence that mental content is externally determined.⁵⁴ I will tentatively accept this conclusion for the purposes of this chapter.

But perhaps Putnam/Burge-style thought experiments that invoke Natural Kinds are not necessary to show the appeal of content externalism. It seems obvious that we can think about the world. A more direct route to externalism, therefore, may come from thinking about proper names. Take the name Prime Minister Trudeau, as I use it. If I am asked for an explanation of its significance I would quite simply point to the fifteenth Prime Minister of Canada and say it is his name. He is the thing that has been assigned the name Trudeau. I take it that it cannot have been assigned to him in a world where he never existed.

Suppose on Twin Earth there was a Twin Trudeau, whom they called 'Trudeau'. Even if superficially like Trudeau, obviously the name 'Trudeau' in that world would mean Twin Trudeau, not Trudeau. So although being remarkably similar to Trudeau (Oscar), Twin Earthians are not speaking the same language or thinking about the same thing. If this is correct, then at least some content is externally individuated. Again, I will tentatively accept this conclusion for the purposes of this chapter since my aim is not defend content externalism, but to merely convey its general appeal, and then assess its compatibility with the kind of epistemic internalism I am defending.

Before moving on to epistemic internalism/externalism, I want to raise one puzzling issue about content externalism. Content externalists are committed to saying that in order to have a thought about something, it is necessary that one have a concept of

⁵⁴ This is far from conclusive, however. For example, surely the meaning of 'I' is internally determined. If Oscar and his twin each think a thought employing the first-person indexical, the thoughts have different truth conditions – even though their meaning is determined internally.

that thing. To have a concept of a thing, they maintain, it is necessary that such things exist in one's environment. Given this, it is not entirely clear to me what a content externalist can plausibly say about thoughts we have about non-existent objects. Presumably, if unicorns existed, they would be a natural kind, which, through either direct or indirect contact with unicorns, I could acquire the unicorn concept and thereby have thoughts about unicorns.⁵⁵ But unicorns do not exist, so, given externalism, I do not have the unicorn concept? Consequently, if externalism is true, can I not have unicorn thoughts?

For example, if this interpretation is correct, I cannot believe that if unicorns existed, they would have a horn; or for that matter, I cannot even believe that unicorns do not exist. It seems counter-intuitive to say that when I believe that unicorns do not exist I do not believe anything at all. I do – it is that they do not exist! I am not sure how externalists propose to answer satisfyingly this query. I raise it in part to cast doubt on externalism since if content externalism is false, it is not incompatible with epistemic internalism. But I also raise the concern to provide an opportunity to clarify what the externalist commitments are.

In response I suppose that a committed externalist could say that it is obvious that we do think and talk about external objects, and our claims and thoughts do relate to them. They could go on to admit that there are some other cases, such as thought about unicorns and other non-existent objects, to which their general account does not apply,

⁵⁵ On the other hand, perhaps 'unicorn' is a mock natural kind term. There are no unicorns to observe so how can we say what they are like (other than that they do not exist)? We do have a conception of their superficial features such as their shape and colour. But superficial features do not define a species. In fact, we all allow that there could be two animal kinds that are superficially identical, and yet be different species. If so, this might suggest that 'unicorn' is an ersatz natural kind term, in which case the term may be semantically deficient, which in turn might make thoughts about unicorns a special case of what is true of thoughts in general.

and these therefore pose a special problem. The fact that there are such special cases does not mean that the externalist treatment of the normal case is incorrect. Or so they could maintain. Also, in anyone's view we can construct concepts. To take Hume's example, we have the concept of 'the Golden Mountain', even though the concept is not instantiated. Instead, it seems to be built up of the more simple concepts 'gold' and 'mountain'. Again, I will have to leave these issues unresolved here, but just raise these issues in part to cast doubt on externalism since if content externalism is false, it is not incompatible with epistemic internalism, and also in part to provide an opportunity to clarify what the externalist commitments are.

With this potential problem highlighted, I will now turn to Timothy Williamson's claim that epistemic internalism is incompatible with content externalism.

3.4. A Reply to Williamson on the Incompatibility of Epistemic Internalism and Content Externalism:

On Williamson's construal of content externalism, "two internal duplicates may differ in what they believe or have propositional attitudes towards." (Williamson, 2007, p. 2) He takes this conclusion to be established by the Putnam-Burge style thought experiments I introduced above. Williamson's expresses the basis of incompatibility as follows:

It is not in dispute that we can pick an example in which Oscar's belief that there are pools of water is justified. Perhaps he is swimming in one. Thus Oscar has the justified belief that there are pools of water. But Twin Oscar lacks a justified belief that there are pools of water, because he lacks the belief that there are pools of water. Thus Oscar and Twin Oscar differ in their justified beliefs, even though they are internal duplicates. Likewise, of course, Twin Oscar has the justified belief that there are pools of twater, while Oscar lacks a justified belief that there are pools of twater, because he lacks the belief that there are pools of twater: that is just another difference in justified belief between Oscar and Twin Oscar. (Williamson, 2007, p. 3)

Williamson's point seems to be that Oscar and his Twin are internally alike and yet, given externalism, there are differences in their justified beliefs, which according to him is contrary to the epistemic internalist's thesis. So if the two are internally identical, then what they are justified in believing must be the same. But since what they are justified in believing is different, epistemic internalism is false since it maintains sameness of justification between counterparts.

The question is, are Oscar and his twin internally identical? It seems so. They are internally the same, i.e. crucially for our purposes there is nothing subjectively distinguishable between them, but the content of their thoughts is nevertheless not internally determined. What move can be made against Williamson's charge that

epistemic internalism is incompatible with content externalism? As we will see, everything depends upon what is meant by ‘sameness’ of justification.

The first thing to note is that the New Evil Demon Intuition supports the awareness condition, that is, one’s counterpart is equally justified in believing the content he does on the basis he does, as we are in believing our counterpart contents based on counterpart evidence in our world. For example, where we have a justified belief about water based on awareness of water, our twin is justified in believing things about t-water on the basis of his awareness of t-water. To be sure, our twin has *no more or no less* justification than we do for believing what he does, nor is he justified in believing no more or no fewer things than we are, for we are both aware of the same kinds of things. In that sense our justification is the “same”, and this sameness is explained in part by the fact that the awareness condition is met. But that is not to say that internal duplicates are *identically* justified, since for one, they are believing different propositions (on the assumption that content externalism is true).

What is also important is that the subjects are internally alike in an important sense. This is NOT in the sense that the content of their belief is the same (since it is not given content externalism), but that if what it is like for them is the same, that is, if there is no phenomenal difference between them, then the same kinds of propositions are justifiable for them, whether or not they believe them.⁵⁶ The sameness in how things seem to them must mean that any differences between them must be ones that transcend

⁵⁶ This reflects the distinction between propositional and doxastic justification. One way of putting this distinction is that propositional justification is what one has justification to believe. Doxastic justification, on the other hand, is when one is justified in believing as one does, *based* on the grounds one does. Doxastic justification, therefore, is propositional justification plus proper basing. This is the difference between believing that *p* with justification as opposed to simply *having* justification to believe that *p* but either: do not believe that *p*, or, one believes that *p* but for the wrong reasons. See chapter 1 as well as Audi’s *The Structure of Justification* (Cambridge: Cambridge University Press, 1993) chs. 7 and 9 for his precise formulation of the distinction.

what they are in a position to be conscious of; what is accessible to them must be the same in the sense of being indistinguishable to them if they consciously attend to it.

What I am suggesting is that a phenomenally indistinguishable experience can be shared by Oscar and Twin Oscar and be the basis of their justification for different propositions. The New Evil Demon case, therefore, also supports Fallibilism, conceived here as the view that a subject can be justified in believing something false, since the New Evil Demon case can be set up to provide a vivid example of subjects who believe falsely, but are nonetheless justified. Suppose Oscar and his twin are being ‘appeared to water-ly’, that is, they have perceptual experiences that are phenomenally indistinguishable from each other of a clear, colourless, odourless liquid that fills lakes and streams. In fact, this stuff is phenomenally indistinguishable from what on Oscar’s planet is *water* and what is on Twin Oscar’s planet is *twater*. If one were being appeared to by water, one by twater, neither could tell the difference if presented with the other. Suppose further that the liquid they are actually presented with is neither water nor twater, but rather a nearly identical liquid, save for the fact it is composed of the molecules QRS. Now if content externalism is true, neither Oscar nor his twin can have thoughts about QRS since neither possesses the concept.

But what beliefs does the perceptual experience of QRS justify? It justifies Oscar and his twin, for example, in believing the following propositions: that the substance appearing before him is drinkable; that the substance appearing before him is wet; that the substance appearing before him is the one that fills lakes and rivers, etc. Importantly, such an experience would justify Oscar in believing that he is being presented with *water*; similarly (though not identically), Twin Oscar is justified in believing that he is being

presented with *twater*. In effect, given epistemic internalism, the same phenomenal visual experience can justify contrary propositions, even if the intentional content of that experience is a different proposition, making what they believe false. Hence justified belief is fallible.

Furthermore, it seems Williamson distorts what it means to say that if two people are internally identical, then are justified in believing the same thing. I take this NOT to mean that if two people are internally identical then they necessarily must *believe* the same justified propositions; but rather, if two people are internally identical, then they are necessarily *justificationally* identical. That is to say, there is a sameness of justification, not sameness of content. The environment may determine which content a subject believes, but it does not determine which contents the subject is justified in believing. That is to say, if two people are internally alike in the relevant respects, then if one has a justified belief and one's twin has the *same* belief, it will be justified and justified to the same extent as the first. A justificatory difference would be if Oscar has the justified belief that *p*, and Twin Oscar had an unjustified belief that *p* where *p* is the *same* content in both cases. Twin Oscar merely lacking the *belief* (justified or otherwise) that *p* because he is unable to form it, given his concepts and his environment, is not a justificationally relevant difference. Rather, that is just difference of content – which is exactly what content externalism requires.

In summary, I have argued that although externalism about content may well be true, internal similarities, in the sense of internal to the subject's conscious awareness, do determine justificational similarities. These similarities are not at the level of what propositions the subject is justified in believing, but at the level of the structure of

justification. That is, internal duplicates are justified in believing exactly the same number of propositions and to exactly the same extent, on the basis of the same internally accessible grounds. Further, the propositions held by the counterpart subjects will themselves be counterparts, e.g. they are as similar as possible while still being non-identical, as well as play an identical role in the subject's noetic structure.

Given this, as well as what I have argued above, content externalism is compatible with the awareness requirement, a necessary condition on justification, and the defining feature of internalist justification.

3.5. Other Reasons to Reject the New Evil Demon Intuition: McDowell's Epistemological Disjunctivism:

In part 3 of “Criteria, Defeasibility, and Knowledge”, John McDowell argues for a form of epistemological disjunctivism, that is, a view that holds that two indistinguishable perceptual experiences can differ in that they put the subject in two different positions to know two different things.⁵⁷ In the experience where there is a fact being ‘made manifest’ to one, the subject is in a position to know that *p*, whereas in the situation where the fact is not made manifest, the subject is merely in a position to know that it merely seems to him that *p*. In developing this view McDowell considers and rejects a version of the Argument from Illusion, the standard argument used to try to show that indistinguishable perceptual experiences are of the same epistemic value or of the same fundamental kind. I think that even if McDowell is right that the argument is not sound and that two indistinguishable perceptions can differ in what they put the subject in a position to *know*, this does not affect what the subject is *justified* in believing.

So while McDowell may be right that the subject's epistemic status is not identical in the veridical and in the hallucinatory case (since they differ in knowledge), it does not follow that they have different *justificatory* statuses. My theory is committed to the claim that the two cases are the same in respect of justification, and this is supported by the fact that we judge subjects and their demon-deceived counterparts to enjoy sameness of justification (in the sense I explained above). I take this intuitive judgement

⁵⁷ As Paul Snowdon has recently pointed out, McDowell's version of disjunctivism is cast purely epistemologically, that is, the view is about the difference between two experiences cast in terms of what one is in a position to know. This is different from the kind of disjunctivism that he and others defend in the philosophy of perception that holds that i) the case of veridical perception is a fundamentally different kind of state than a hallucinatory state and ii) that it does not follow from the fact that two experiences are indistinguishable that they are of the same kind. According to McDowell's disjunctivism the dispute is epistemological, whereas the one Snowdon represents is metaphysical. For details see Snowdon (2005).

to be a datum that should only be rejected if we are given a good reason to do so. So far we have seen that content externalism provides no such reason. We will now consider whether McDowell provides sufficient reason to overturn this judgement. To begin to do so, it is necessary to turn to the version of the Argument From Illusion that McDowell presents.

All sides should agree that our ability to tell how things are by looking at them is fallible – there are cases that deceive us. These deceptive cases are cases where $\sim p$ is the case, although this perceptual experience is qualitatively identical to non-deceptive cases in which p is the case. Advocates of this argument maintain that if one's experience falls short of knowing the fact in the deceptive case, it also falls short in the non-deceptive case as well. The 'highest common factor' shared between these two cases is the one compatible with $\sim p$. As McDowell puts it, the highest common factor between the two cases is "at best a defeasible ground for the knowledge". (McDowell, 1998a, p. 386) The picture is that of these common appearances constituting a 'veil of ideas' that mediate between mind and world.

This might be understood in at least two ways: one is the traditional idea of a literal veil made of ideas. Such 'sense data', which is what we are directly aware of, either *constitutes* the world (idealism), or else *mediates* between us and the mind-independent world, thereby hiding the world from us. The other idea is simply that all we can tell is how things appear to us. The argument that McDowell sketches merely gets us to the second idea and not the first, since it does not follow from the fact that one has perceptual experiences that are subjectively indistinguishable from each other that the objects of perception are mere sense data. At least in the case of illusion, it could be that

the objects of perception are mind-independent objects, but share an appearance. But with either model in place, it is easy to see how one might be led to sceptical conclusions about the extent of our knowledge. This follows on the plausible assumption that it is a necessary condition of knowing that p that one must at least be able to discriminate p from $\sim p$.⁵⁸ By hypothesis this condition is impossible to meet if one's grounds for believing that p are qualitatively identical to one's grounds when $\sim p$ is the case. McDowell sets out to combat such scepticism.

McDowell's response to the argument is to deny that there is necessarily epistemologically common ground in the deceptive and non-deceptive cases. His suggestion is that "an appearance that such-and-such is the case can be *either* a mere appearance *or* the fact that such-and-such is the case making itself perceptually manifest to someone." (Ibid., p. 386-387) One might think that McDowell's disjunctive suggestion looks promising – that in the deceptive and non-deceptive cases the difference

⁵⁸ Principles that have been proposed to capture the relevant notion of 'discrimination' required for knowledge include *Sensitivity* and *Safety*. *Sensitivity* holds that: "If an agent knows a proposition, then that agent's true belief must be *sensitive* in the sense that, had that proposition been false, she would have not believed it." (Pritchard, 2006, p. 182) See Nozick (1981) for an early defense of the 'truth-tracking' theory of knowledge that employs the sensitivity principle.

Safety, by contrast, holds that: "If an agent knows a proposition, then that agent's true belief in that proposition must be *safe* in the sense that it couldn't have easily been false (alternatively: were the agent to continue believing that proposition in similar circumstances, then the belief would almost always still be true)." (Pritchard, 2006, p. 180-1) Another way of putting this is that a subject's belief is safe just in case a subject believes that p *because* p is the case and the belief is stable or robust in the sense that in all of the nearby possible worlds in which the subject would continue to believe that p , p is true in at least most of those worlds.

What both of these principles have in common is that they are epistemically *externalist* in nature. One salient way in which they differ is that *Safety* but not *Sensitivity* allows subjects to know the denials of sceptical hypothesis. A further difference is that *Safety* is consistent with knowledge being closed under known entailment, whereas *Sensitivity* requires the rejection of the highly intuitively closure principle, which seems to count decisively against *Sensitivity*. Closure here is being understood, roughly, as the principle that if one knows a proposition p , and knows that it entails another proposition q , then one knows q as well. This principle is so overwhelmingly intuitive and as close to self-evident as any epistemic principle can be that it should be taken as a datum. As such, if any theory is incompatible with closure, so much the worse for the theory.

As my aim here is to elucidate the nature of epistemic justification, rather than knowledge, and given that for methodological purposes I am presupposing the falsity of scepticism, I will remain neutral on how best to characterize the notions of discrimination and anti-luck that seem to be required for knowledge.

may well be a difference in what one is acquainted with in experience, viz., whether it is a fact or a mere appearance. This may have epistemological implications insofar as they are reflected in differences in what the subject is in a position to *know*. However, there are many different kinds of epistemic evaluation other than knowledge e.g. “Foley-rationality”, or epistemic blamelessness/responsibility, or reasonableness, and most centrally, epistemic justification.⁵⁹ This is worth noting because it may turn out that certain kinds of internalism or externalism may hold for some epistemic concepts, but not for all, and so generalizations across epistemology should be made with great care (see chapter 1).

What is clear is that even if one adopts McDowell’s disjunctive stance, nothing straightforwardly follows to suggest that the deceptive and non-deceptive case are not identical in justification. A complication may be thought to arise, however, due to the fact that McDowell repeatedly stresses that “knowledge – at least as enjoyed by rational animals – is a certain sort of standing in the space of reasons.”⁶⁰ (McDowell, 1998c, p. 395) While exactly what McDowell means by this phrase is the matter of some debate⁶¹, one would assume that we have some prior grasp of what constitutes a reason, e.g. a consideration that counts in favour of something. The idea seems to be that whether or not one knows depends on things to do with reasons. Whatever else may be true of reasons, it seems that they are related to justifications. For example, it seems plausible to suppose that reasons are those things that justify beliefs.

⁵⁹ By traditional lights, justification is necessary but not sufficient for knowledge. But I remain agnostic on this point. That is not to say, however, that justification is not an interesting and important concept of epistemic appraisal in its own right.

⁶⁰ This motif of the ‘space of reasons’, which McDowell attributes to Wilfrid Sellars, is repeated frequently throughout his work. See especially McDowell (1996).

⁶¹ See for example essays in Smith (ed.) (2002).

If this is what McDowell means by reasons and their space, then I hold that McDowell is conceiving of justifications and reasons in an inaccurate way. This is because he holds that in a case of veridical perception and a case of hallucination we are differently placed vis a vis knowledge, whereas the subjects are similarly placed in relation to justification/reasons, as the case of the new Evil Demon reveals. So McDowell is making some mistake about reasons: either they are not required for knowledge, and therefore knowledge is not a standing in the space of reasons; or else reasons are required for knowledge, and since the same reasons/justification is present in both the veridical and hallucinatory case, the fact that the subject fails to know in the hallucinatory / illusory case is due to some extra-justificational feature of the situation, e.g. that what he believes is false. A further consequence of this is that McDowell is mistaken that there are factive empirical reasons, on the assumption that it is reasons that justify belief, given that the same reasons are present both in the veridical and the hallucinatory case. I shall add more on the coherence of the notion of factive empirical reasons below.

It is important to note that there is a 'common factor' between the deceptive and non-deceptive cases, even if this common factor is not a metaphysical one at the level of perceptual states (one's mental state) or a common state of knowledge. What is common is a phenomenally indistinguishable appearance. In the context of discussing the so-called Phenomenological Argument for the Highest Common Factor View, McDowell himself concedes: "The alternative conception [McDowell's brand of disjunctivism] can allow what is given to experience in two sorts of case to be the same *in so far as* it is an appearance that things are thus and so; that leaves it open that whereas in one kind of

case what is given to experience is a mere appearance, in the other it is the fact itself made manifest.” (emphasis in original) (Ibid., p. 389)

This sameness of appearance that things are thus and so is what I argue is the basis of epistemic justification which must be so, given that a subject and her deceived twin are both justified and have all the same things accessible to them. The reason that the appearance of a fact making itself manifest is not justificationaly relevant is because if it were, there would be a difference in justificatory status between a subject and her twin; but as careful consideration of the case reveals, there is no such difference. When a fact makes it itself manifest, what the subject has access to is the *fact* itself (that is, an appearance of a certain kind), not the fact *that* it is a fact making itself manifest. From the subject’s own perspective, what is accessible in the deceptive case is exactly what is accessible in the non-deceptive case, viz., an appearance that things are thus and so (regardless of what *kind* of perceptual state the subject is in, or what she may be in a position to know). In short, what the subject knows on McDowell’s account is a fact, whereas what the internalist about justification requires is awareness of the ground of his belief, as a ground. What the New Evil Demon case shows is that it is justificationaly irrelevant whether what the subject is aware of is genuinely a fact making itself manifest or a mere appearance – the subject is equally justified in believing as he does, on the basis he does, in both cases.

The above remarks raise the question about epistemic access. I argued in chapter 2 that the kind of access that is necessary for epistemic justification is ‘strong’ conscious awareness, and in the case of immediate justification, *non-doxastic* strong awareness.

Which kinds of things need to be accessible if justification is to obtain? Consider these three different kinds of internalism that Duncan Pritchard and Ram Neta formulate:

- 1) S's epistemic justification for believing that *p* is constituted solely by S's mental states.
- 2) S's epistemic justification for believing that *p* is constituted solely by facts that S can know by reflection alone.
- 3) S's epistemic justification for believing that *p* is constituted solely by properties that S has in common with her recently envatted physical duplicate. (Neta and Pritchard, 2007, p. 382)

In seeing where McDowell fits into the above categories, we can come to see the strengths and weaknesses of each.

The first version of internalism is what I have been calling 'Mentalism'.⁶² As I argued in chapter 2, while Mentalism may capture a general truth, it is misleading as a theory of epistemic justification, since it does not include an essential awareness condition, and as such it should be rejected in the above form (See 2.2). Given that McDowell stresses that knowledge is nothing more than an appropriate standing in the space of reasons, and given that he conceives of reasons as mental states, he also accepts the spirit of 1). Support that McDowell holds that knowledge is nothing more than a proper standing in the space of reasons can be found in the following passage in which he criticizes what he calls the "hybrid conception of knowledge":⁶³

In the hybrid conception, a satisfactory standing in the space of reasons is only part of what knowledge is; truth is an extra requirement. So two subjects can be alike in respect of the

⁶² First introduced and defended, to the best of my knowledge, by Feldman and Conee in "Internalism Defended" reprinted in *Epistemology: Internalism and Externalism*, ed. Kornblith, Oxford: Blackwell, (2001). For more on Mentalism see 2.2.

⁶³ The 'hybrid' conception of knowledge, I take it, is one that conceives of knowledge in terms of 'inner' and 'outer' features; i.e. a picture of knowledge as a relation that obtains between a believer who possesses good reasons (which are internal to the subject) and a fact or a true proposition (which is external to the subject).

satisfactoriness of their standing in the space of reasons, although only one of them is a knower, because only in her case is what she takes to be so actually so. But if its being so is external to her operations in the space of reasons, how can it not be outside the reach of her rational powers? And if it is outside the reach of her rational powers, how can its being so be the crucial element in an intelligible conception of her knowing that it is so – what makes the relevant difference between her and the other subject? Its being so is conceived as external to the only thing that is supposed to be epistemologically significant about the knower herself, her satisfactory standing in the space of reasons. (McDowell, 1998c, p. 403)

Support that McDowell thinks of reasons as mental states can be found in the following passage:

Someone who remembers that things are a certain way, like someone who sees that things are a certain way, has an excellent reason for taking it that things are that way; the excellence comes out in the fact that from the premise that one remembers that things are this and so, as from the premise that one sees that things are thus and so, it follows that things *are* thus and so. The epistemic positions themselves put their occupants in possession of reasons for their beliefs; those reasons do not need to be supplemented with less cogent arguments from non-question-beggingly available premises. (emphasis in original) (McDowell, 1998b, p. 427-428)

So according to McDowell, one's reason to believe that a desk is in front of one, for example, is one's seeing that there is a desk in front of one, which is a mental state of the subject.

However, for reasons similar to the ones I gave in chapter 2, McDowell should reject this formulation of Mentalism since it includes no awareness requirement. Support for the claim that McDowell thinks that some form of awareness requirement is needed is the following: "I agree with [Elizabeth Fricker] that we lose the point of invoking the space of reasons if we allow someone to possess a justification even if it outside his reflective reach." (McDowell, 1998b, p. 418) On Neta and Pritchard's reasonable

enough interpretation, McDowell uses 'being inside of one's reflective reach' to mean that the agent is able to know it by reflection alone. (Neta and Pritchard, 2007, p. 384)

The reason that McDowell should reject Mentalism as formulated if he accepts a reflective access requirement is that it is not at all obvious that one can always tell by reflection alone what mental states one possesses. Even if one has privileged access to one's mental states, that is not to say that one has infallible access. If one needs reflective access to one's justificatory grounds and this access is fallible or limited, and one's grounds are one's mental states whether or not one has reflective access to them, then it seems that one will have some justifications that one does not have access to, i.e. reasons that are 'outside the subject's reflective reach'. This, McDowell claims, defeats the point of 'invoking the space of reasons', and as such is unacceptable. To put it in my own terms, Mentalism is unacceptable since it violates the Subject's Perspective Objection (SPO), which recall is the following:

If the subject holding a belief is not aware of the belief's grounds, then she is not aware of how its status is any different from a stray hunch/arbitrary conviction/strong but groundless conviction. As such, if the subject regards her belief as justified, she ought to regard the fact that the belief is justified as a brute fact. From that we may conclude that from her own perspective she ought to regard it as an accident that her belief is true, i.e. from her perspective a belief whose justification appears to be a brute fact could just as easily be false. And that implies that the belief is not justified. (See Chapter 2 for more details)

The second kind of internalism, a kind of 'access internalism', is the one that McDowell seems to subscribe to, given his remarks about the point of invoking the space of reasons. He thinks that one's positive epistemic status, in the case of perceptual beliefs at any rate, is determined by facts making themselves manifest to one in one's experience, which if they are to play the role of reasons, must be within the subject's

reflective reach. Recall that McDowell unequivocally states that: “I agree with [Elizabeth Fricker] that we lose the point of invoking the space of reasons if we allow someone to possess a justification even if it outside his reflective reach.” (McDowell, 1998b, p. 418) Again, while it is not entirely clear what being inside one’s reflective reach amounts to, I will follow Neta and Pritchard in understanding it as meaning that a reason is within one’s reflective reach when it can be known by reflection alone.

It is somewhat unclear if what is supposed to be accessible upon reflection is the experienced fact, or the fact that it is a fact. That is, is what is accessible upon reflection, for example, that one has hands when one sees that one does, or that one is *seeing* that one has hands (which entails that one has hands)? If the former, it is not clear if this will count as ‘access internalism’ as Neta and Pritchard construe it; if the latter, it does not seem that one can tell by reflection if, for example, one is really seeing one’s hands, rather than merely seeming to see. More on this below.

It might seem puzzling why McDowell stresses the importance of *reflective* access in connection with space of reasons in the passages I quote above. What is the significance of reflection? Nothing essential, I submit. It is presumably myopia that has led to overlooking central cases of justified believing, such as having beliefs about what one is presently experiencing in sense perception. In that case one does not introspect the world around one, or access it through *reflection*; rather one *perceives* it. What is common to the kind of perception and reflection or introspection that the internalist is moved by, and that McDowell seems to be interested in here in connection to his conception of the space of reasons, is *conscious awareness*. In sum, it does not seem like there is anything particularly significant for McDowell about reflection *per se*. What is

justificationally important about reflection, for McDowell as well as for me, is that it is our way of becoming conscious of the contents of our minds, just as perception is our way of becoming conscious of the world.

Returning to the issue of knowing in the space of reasons, although McDowell does not explicitly say that having facts making themselves manifest to one in one's experience is sufficient for justification, since he does not say any more either way (apart from requiring the absence of defeaters), and claims that justification is *solely* a function of what is reflectively accessible, this leads me to conclude that he is so committed.

One problem with this kind of internalism arises if it is offered as a theory of doxastic justification. It is possible to experience a fact, but not base any relevant belief on it. I contend that the basing relation is necessary (but not sufficient) for *doxastic* epistemic justification. Any viable theory of justification needs to respect the epistemic difference between believing that *p* and having a good reason to believe it, and believing that *p because* of the ground, that is, on the basis of the ground. A fact made manifest, and hence 'internal' in McDowell's sense, is insufficient for justification. But crucially for my purposes, the New Evil Demon case shows that a fact being made manifest is not necessary for justification either, since a subject can be equally justified in believing something if it is based on its *seeming* to the subject that *p*, rather than *p* itself really being the case. But perhaps McDowell is offering his position as a theory of propositional justification, in which case the above criticism would not apply.

But there is an even more fundamental objection: is McDowell's version of internalism even coherent as Pritchard and Neta suggest? The unique McDowellian

thesis is the conjunction of two claims, which I will follow Pritchard and Neta as understanding as follows:

Reflective Accessibility of Reasons: (RAR) One is in a position to know, by using one's reflective capacities, what one's reasons are for believing that *p*.

Factivity of Reasons: (FAR) There are factive empirical reasons for beliefs about the external world. (Neta and Pritchard, 2007, p. 388-389)

The conjunction of these two theses gives us:

(MCD) One is in a position to know, by using only one's reflective capacities, what one's reasons are for believing that *p* – even when those reasons are (as they sometime are) factive empirical reasons. (Ibid. p. 389)

Whether or not this view is coherent depends on not only the coherence of RAR and FAR, but also on their conjunction. Is it really possible to be in a position to know what one's reasons are if they are factive? Does this not require knowing *that* they are factive, and this is not something that is knowable upon reflection?

For example, if one sees that *p*, one can 'access one's reasons' in the weak sense of knowing that one is having a perceptual experience with a certain phenomenal character, but one cannot tell by introspection alone whether one is seeing that *p*, or merely *seeming* to see that *p* (i.e. seeing a situation in which *p*). That is, what content a state has is not something that can be read off how things appear to the subject. For example, seeing that the cup is chipped and *seeming* to see that the cup is chipped are different states with different truth-conditions, and yet they appear identically to the subject who experiences them.

If according to McDowell one's reason to believe that *p* is the factive state of one's seeing that *p* (and not the non-factive but phenomenologically indistinguishable state of one seeming to see that *p*), then 'knowing what one's reasons are' seems to require being able to discriminate one's factive perceptual state from one's qualitatively

identical non-factive state. But this is impossible, since there is nothing subjectively to distinguish between them. In short, it is false that a subject can know what his reasons for a belief are if those reasons are factive. Knowing the facts are not enough to know them as the reasons that justify one's belief, if what is justificationaly relevant is their factiveness, since this is something that transcends the subject's reflective awareness. If what is key is that only what can be known by reflection alone can be justificationaly relevant, then reasons for empirical belief about the external world cannot take a factive form.

Notice that the non-factive conception of reasons that my account of epistemic justification presupposes does not face this incoherence, since in effect I endorse RAR but not FAR. In paradigm cases of justified belief, which therefore meet the Awareness Condition, the subject is in a position to know, by using her reflective capacities, what her reasons are for believing as she does. A problem here about knowing what one's reasons are by reflection alone might arise in the case of doxastic (rather than propositional) justification, depending on the nature of the basing relation. That is, a problem may arise if the relation between one's belief and one's reason when one's belief is held *for* that reason is a causal relation of some kind. This would be a problem because knowing what causes what is not an *a priori* matter.

However, at least in the case of *propositional* justification, the basing relation plays no role, in which case one can know by reflection alone what one's reason is for holding a particular belief.⁶⁴ For example, when one believes that one has hands on the basis of seeming to see hands at the end of one's arms, one is aware that one's reason for

⁶⁴ See 1.3. for a discussion of the basing relation. See chapter 2 for an explanation and defense of the notion of non-doxastic awareness.

this belief is one's perceptual experience as of hands. One is aware of how, or in virtue of what, one is justified in believing as one does. While there are empirical reasons for beliefs about the external world, which are factive as FAR maintains, they are not reasons *in virtue* of their factive nature – factive and non-factive states with the same phenomenal features are equally good (or bad) reasons that can justify empirical beliefs about the external world.

If Mentalism and the standard form of access internalism are inadequate theories of epistemic justification, the best remaining candidate is Neta and Pritchard's 3): S's epistemic justification for believing that *p* is constituted by properties that S has in common with her recently envatted physical duplicate, which must be so to account for our evaluative judgement that a subject and her duplicate in vat do not differ in what they are justified in believing.

The principle about envatted replicas follows from our evaluative thesis about sameness of justification between counterparts. But is our claim about sameness of justification not derivative from something more basic? It is difficult to say. But what is important here is that our claim to sameness of justification is an *evaluative* thesis, not an *explanatory thesis*.⁶⁵ It says only that twins internally alike are justificationaly alike; not *why* or in virtue of what are they so justified. As Nico Silins has recently pointed out, this is a virtue of evaluative theses in general, since what supports them is pure intuitive plausibility, not motivations from theoretical judgements about the nature of justification. (Silins, 2005, p. 386) As relatively pre-theoretical judgments, they can be used as data in philosophical theory building, as I have been doing here. Such data should be used

⁶⁵ See Silins (2005) p. 385 for discussion of evaluative v. explanatory theses in epistemology.

unless we have good reason to reject it, which is of course what I have been evaluating throughout this chapter.

The common properties between internal twins are that things appear to the subjects to be a certain way. This is the case even if metaphysical disjunctivism about perception is correct and these two experiences are in fact of two different ontological kinds. Similarly this thesis remains intact even if McDowell's epistemological disjunctivism is broadly correct and the subjects are in positions to know different things, although the details of the position will need to be different since the conjunction of RAR and FAR is of dubious coherence, as we have seen. Still, perhaps there is *some* kind of reflective access that can be conjoined with McDowell's thesis that there are factive empirical reasons, which if true, might be relevant for knowledge since an agent can know things in the veridical situation that his recently envatted counterpart could not know. But such an approach does not apply to non-factive modes of epistemic evaluation, such as justification traditionally construed, which does not vary unless there is a change in how things seem to the agent.

In summary, I have argued that we can have unreliable but justified beliefs, as cases involving the New Evil Demon show. The case of the New Evil Demon shows *sameness* of justification between counterparts in a demon and a non-demon world. What this shows is that since the subject is justified but *actually* unreliable in the demon world, then he is justified in the non-demon world whether he is actually reliable or not, so long as the properties he shares with his demon world counterpart are still instantiated. That the subject's belief-forming mechanisms happen to be reliable in the non-demon world is beside the point vis-à-vis his epistemic justification. The case of the New Evil

Demon also supports Fallibilism, understood here as the position that one can be justified in believing something that is false. As we have seen, neither arguments regarding content externalism nor McDowell's version of epistemological disjunctivism count against these conclusions.

For all I have shown above, *actual* reliability or truth-conduciveness is not necessary for epistemic justification to obtain, but it could still be the case that the reason that the subject is justified in the demon world is precisely because his beliefs *would be* reliable in the non-demon world. So even if *actual* reliability is not required, might *conditional* reliability be necessary for epistemic justification to obtain? If so, then there would be a (external) condition on which epistemic justification depends that would not meet the Awareness condition, in which case epistemic internalism as I construe it would be false. It is to these issues that I now turn.

CHAPTER FOUR

THE TRUTH-CONNECTION: WHAT IT IS NOT – RELIABILITY REDUX

4.1. Introduction:

As I noted in Chapter 1, one of the central characteristics of justification is, as Alston says, that it has “to do with a specifically *epistemic* dimension of evaluation.” (Alston, 1989a, p. 83) As I noted there, the use of the qualifier “epistemic” is needed to individuate epistemic justification from other possible sorts, such as moral, aesthetic, or pragmatic justification. I agree with Alston that *some* connection to truth is what makes a given thing epistemic. (See 1.2. for the four basic features of justification) In this chapter I will investigate proposals that conceive of the relation between justified belief and truth as one of truth-conduciveness. I will argue that, given internalism about justification, either in the sense I defend in Chapter 2 involving the Awareness Requirement, or in a weaker sense that Christopher Peacocke endorses, the truth-relation cannot be characterized in terms of truth-conduciveness or reliability. This is not just to reiterate a conclusion of Chapter 3 that justification need not be *actually* truth-conducive. I will argue further that the truth-connection cannot even be one of truth-conduciveness in a *conditional* mode. Having dispatched truth-conducive accounts of the truth-connection, I will be in a position to develop a positive account of the truth-connection in Chapter 5.

Here and throughout I will use ‘truth-conduciveness’ and ‘reliability’ interchangeably to denote the property of a tendency to result in true beliefs. So to say of a belief

that it is likely to be true, therefore, means roughly that the method that gives rise to the belief is likely to deliver true results in the circumstances deployed. This is, of course, quite vague, but the notion of reliability is a notoriously difficult concept to make precise in any non-trivial and informative way. I take this as a further consideration counting against employing the notion of reliability in philosophical theorizing. I will return to this issue of how to characterize reliability at the end of this chapter.

To appreciate the *prima facie* threat to a substantial truth-connection, which would undermine the notion that the account of justification I have been defending is genuinely epistemic, recall that I have argued that actual truth-conduciveness or reliability is neither necessary nor sufficient for epistemic justification.⁶⁶ In Chapter 2 I argued that we can have reliable but unjustified beliefs, as shown by cases of clairvoyance and other cases in which the subject lacks conscious awareness of her grounds. In Chapter 3 I argued that we can have unreliable but justified beliefs, as cases involving the New Evil Demon show.

The case of the New Evil Demon shows *sameness* of justification between counterparts in demon and non-demon worlds. What this shows is that since the subject is justified but *actually* unreliable in the demon world, then he is justified in the non-demon world whether he is actually reliable or not, so long as the properties he shares with his demon world counterpart are still instantiated. (See Chapter 3) That is, since the subject is justified in a demon world, but not reliable, reliability is not necessary for justification. That the subject's belief-forming mechanisms happen to be reliable in the non-demon world is beside the point vis-à-vis his epistemic justification. Given that this

⁶⁶ I take my arguments as also directed against accounts that conceive of the connection between justified belief and truth as one of *entailment*. This is because I am treating truth entailment as the limit case of truth-conduciveness.

is the case, what is the relation between justified belief and truth? For all I have shown in Chapter 3, the truth-connection cannot be one of *actual* reliability or truth-conduciveness, but it could still be the case that the reason that the subject is justified in the demon world is precisely because his beliefs *would be* reliable in the non-demon world. In this chapter I will block this move: reliability in either its actual or conditional mode is not necessary for epistemic justification to obtain.

Given that actual reliability has been ruled out as an account of the truth-connection, I will turn to conditional accounts of reliability, beginning with Alvin Goldman's Normal Worlds Reliabilism (Goldman, 1986). By seeing how this account is *prima facie* incompatible with various forms of epistemic internalism, I will evaluate Christopher Peacocke's recent attempt to reconcile his form of internalism with the kind of conditional reliability account that Goldman introduces. Seeing how Peacocke's account fails will reinforce the conclusions of Chapter 3 that truth-conduciveness (actual or conditional) is not the correct account of the truth-connection for internalists. I shall try to show that an account of the truth-connection in terms of truth-conduciveness is not motivated, but even if it were, Peacocke's account (and others like it) sets an impossible burden to meet, which results in its being inconclusive as to whether or not any of our beliefs are justified. As such, the theory must be rejected, as common sense confidently assures us that the beliefs on which Peacocke's theory casts doubt are indeed justified.

In short, I argued in chapter 3 that *actual* reliability is not a necessary condition for justification; here I will argue that *conditional* reliability is not necessary for justification to obtain, either. Taken together, these chapters aim to show that external

factors are not necessary for justification to obtain, and so epistemic internalism is further supported.⁶⁷

⁶⁷ Even if it is a requirement of internalism that reliability is not necessary for justification, this of course leaves open what the relationship is between justification and knowledge, as well as it if there are other useful epistemic evaluations for which reliability plays an essential role.

4.2. The Subject's Role and the Truth Connection: Normal World Reliabilism

The New Evil Demon cases show that if a subject's belief that p is epistemically justified, the belief that p need not actually be likely to be true – that is, the belief need not actually be reliably produced or sustained. The connection between justified belief and truth must lie elsewhere. At this point one might naturally think that the reason that the subject is justified in the demon world is precisely because his beliefs *would be* reliable in the non-demon world.⁶⁸ And one could adopt this view short of 'mad-dog' externalism⁶⁹, holding that reliability is necessary for justification, but also maintaining that what the subject is (or easily could be) aware of is also necessary for justification. Such an approach tries to accommodate a role for reliability, while holding that the subject's point of view is also essential to epistemic justification.

Alvin Goldman's Normal World Reliabilism seems to be an account that could be brought into service to support such a view. Goldman's theory says that a belief is justified if and only if it is likely to be true in 'normal worlds'. That is, the belief need not be likely to be true in the actual world. Instead, reliability is relativized to so-called 'normal worlds'. Normal worlds are fixed by what beliefs the subject holds, and so

⁶⁸ Here and throughout I have been speaking of demon worlds, meaning a world of a kind that Descartes considers in his first meditation. But nothing essential rests upon choosing this sceptical scenario, as such a hypothesis is meant to be merely representative. Other skeptical scenarios besides the Cartesian evil demon include being an envatted brain at the hands of evil super-scientists, being asleep and having a lucid dream, etc. These scenarios will be taken as sceptically equivalent for present purposes, and accordingly, are intersubstitutable *salva veritate*.

For short, all these sceptical scenarios can be collectively referred to as "philosophical dreaming". Sosa defines philosophical dreaming thus: "S philosophically dreams p iff S experiences as if p , but *unveridically* so." (Sosa, 2003, p. 143) So in the text when I refer to a non-demon world, I at least mean one where the evil demon hypothesis is false, and, that by and large, one's perceptual experiences are veridical.

⁶⁹ 'Mad-dog externalism' is understood here as the view that the subject's perspective is justificationaly irrelevant, and that factors external to it, e.g. the reliability of the process that gave rise to the belief in question, are necessary and sufficient for epistemic justification to obtain.

accommodates the intuition that the subject's perspective necessarily affects the truth-connection. Goldman writes,

We have a large set of common beliefs about the actual world: general beliefs about the sorts of objects, events, and changes that occur in it. We have beliefs about the kinds of things that, realistically, can and do happen. Our beliefs on this score generate what I shall call the set of *normal worlds*. These are worlds consistent with our *general* beliefs about the actual world. (original emphasis) (Goldman, 1986, p. 107)

So the idea is that while justified beliefs may not be likely to be true in the actual world (say, because the actual world is a demon world), they *would* be true if the actual world were among the set of 'normal worlds', where these are worlds where most of our general beliefs are true. We can put this idea like this:

Conditional Reliability Version of the Truth Connection: If S's belief that *p* is justified, then *p* is likely to be true in worlds where S's other general beliefs about the way the world is are true.

I will say more at the end of this chapter about difficulties in trying to make the notion of 'normal worlds' precise.

As stated, the conditional version of the truth connection is a form of epistemic externalism, since the reliability of a belief-forming mechanism and what world it is being used in are not things that the subject can be aware of on reflection alone (if at all). However, different *kinds* of internalists and externalists alike could hold that this truth connection is *necessary* for justification; although for reasons I will explain, no such truth connection is even necessary for epistemic justification. Externalists' accounts of justification could hold that the principle as stated is *sufficient*: if a belief is likely to be true in worlds that would make true most of the subject's other beliefs, then the belief is

justified. A more qualified externalism could hold that such a conditioned truth connection is sufficient for justification in the absence of defeaters.

On the other hand, an internalist might hold that a truth connection of this kind is necessary for justification, but the subject also must have some kind of cognitive access to the fact that this relation obtains. The crudest of such approaches would consist of doxastic accounts that Stewart Cohen calls “Intellectualist” theories of justification. He construes Intellectualist theories as follows: “In order for S to be justified in believing B, S must believe that certain conditions obtain which make the truth of B probable.” (Cohen, 1984, p. 286) So, combining the Intellectualist and Conditional Reliability approaches, one gets the following view:

The Intellectualist Conditional Version of the Truth Connection: In order for S’s belief that *p* to be justified, S must *believe* that *p* is likely to be true in worlds where S’s other general beliefs about the way the world is are true.

But an advocate of this kind of position should not be satisfied with requiring the subject to merely believe the conditional that *if* he were in a ‘normal’ world, then his justified beliefs are also likely to be true. What we are interested in is not just whether the subject believes that his beliefs *may* be justified, but whether he has reason to regard them as *actually* justified, that is, does the subject have reason to regard the token beliefs he holds as justified in the world he finds himself in (normal or otherwise)? To establish this conclusion, the subject would also have to believe that the actual world is a ‘normal world’; or if the actual world is not ‘normal’, that how he formed his belief would be reliable if it had been formed in a normal world. And, presumably, it is not enough that the subject *merely* believes that the actual world is a ‘normal world’, or that it is one where how he formed his belief is conditionally reliable, even if it is, because such a

belief could be held for absurd reasons, or for no reason at all. Someone who advocates an 'Intellectualist' model of justification, presumably, would insist that the subject's second order belief about the connection between his first order belief and truth must be *justified*.

At this point one might think it obvious that no such Intellectualist account could possibly be correct on pain of scepticism. First, does the Intellectualist face a vicious regress of increasingly complex justifications? Second, it seems that people just do not tend to have beliefs, justified or otherwise, about the reliability of their belief-forming mechanisms. Cohen himself rejects such accounts for just this reason. (Cohen, 1984, p. 286) But he goes on to note that an Intellectualist could respond to this objection in two ways: he could either maintain that people *do* have the relevant supporting beliefs, but have them subconsciously or that the beliefs are somehow held dispositionally. (Ibid.) The first approach is obviously an empirical matter, and cannot simply be presupposed. Cohen seems to reject the dispositionalist account on the grounds that it is difficult, if not impossible, to determine when a belief is actually held dispositionally, and when it is just something that someone would affirm (for the 'first time', as it were) if it were entertained.⁷⁰ Despite this obvious epistemic difficulty, the distinction between a subject having a disposition to entertain a belief already held, and the subject merely acquiring the belief if he considered the relevant proposition, seems genuine.

A more promising response is not to insist that people do have these beliefs, but that there is justification to believe them. Christopher Peacocke defends a variant of the

⁷⁰ Robert Audi defends a similar distinction between *believing* and *affirming*. For discussion of this point, see Audi (1982).

kind of Intellectualist approach I have been describing. Peacocke in effect takes up the kind of challenge that Alston issues in *The Reliability of Sense Perception*. Alston asks,

Why suppose that sense perception is, by and large, an accurate source of information about the physical environment? [...] More generally, why suppose that any of the bases on which we regularly and unquestionably form beliefs are reliable, can be relied on to yield mostly true beliefs? [...] Do we, or can we, have any solid reason for this assurance? [...] Do we possess, or can we obtain, adequate reasons for supposing that this is the way things are?" (Alston, 1993, p. 1)

Alston ultimately answers this question in the negative, arguing that all positive attempts to answer this question fall prey to vicious forms of objectionable epistemic circularity. (Alston, 1993) Taking a more optimistic line, as I will explain below, Peacocke's account tries to reconcile Intellectualist elements with a conditional reliability version of the truth-connection in trying to answer Alston's challenge. I will call this view *Peacocke's Internalism*, which I will refine below, but which essentially contends that it is possible to *show* or *prove a priori* that we are justified in believing that the conditional reliability relation holds between our justified beliefs and truth.

Peacocke's view is different from the Intellectualist account in that he does not think we all necessarily have beliefs about the reliability of our belief forming mechanisms, but his account preserves Intellectualist elements by maintaining that we can be justified in such beliefs *a priori*. The account is a form of internalism, but not in my sense that we must be aware of the grounds of our justified beliefs. Peacocke's account does share internalist elements in the sense that he does not think that reliability is sufficient for rationality. Further, he holds that it is not just that we can be justified in believing that *if* we are in a normal world, our beliefs are likely to be true. Rather, he thinks that he can show *a priori* that we have reason to think we actually inhabit such

worlds, and that our belief-forming processes are likely to be reliable in them. That is, he thinks he can show that we can be justified in believing that a necessary condition of justification holds.

However, one might still wonder if being able to show in principle *a priori* that we are in a world where belief formation processes are likely to be reliable will ultimately be of much help in defending epistemic internalism. Even if Peacocke's arguments are sound, they will not be of help to the person who is not aware of these arguments, or to the person who, while aware of the existence of such arguments, cannot follow the reasoning. It is true that Peacocke's approach will not address the general first-personal or egocentric concerns that I argue motivate the form of epistemic internalism I am defending. Nevertheless, if sound, Peacocke's arguments would help combat the most radical forms of epistemic externalism, and if not support internalism, at least support a form of non-externalism. Also, to those who were apprised of Peacocke's arguments and could follow them, they would have the resources to be epistemically justified in my internalist sense. Such arguments would allow for a kind of internally accessible rational defensibility for the reliability of sense perception. As such, the compatibility of some form of awareness internalism (though perhaps not in its strongest forms) depends on how successful accounts like Peacocke's incorporate reliability conditions with Intellectualist elements.

Accordingly, I will now turn to the details of what I will call Peacocke's Internalism, and assess its coherence.

4.3. Peacocke On Entitlement and Truth-Conduciveness:

In *The Realm of Reason*, Christopher Peacocke develops a ‘generalized rationalism’ concerning, among other things, what it is for someone to be ‘entitled’ to form a given belief. While Peacocke’s work focuses on what he calls *entitlement* rather than justification, much of what he says could be applied *mutatis mutandis* to the theory of justification, and so is worth evaluating here. This is because Peacocke’s basic conception of entitlement shares many features in common with justification that I outlined in Chapter 1. For instance, as I highlighted in connection with justification, Peacocke similarly maintains the following about entitlement: it applies to beliefs, but perhaps derivatively so, since Peacocke’s focus is as follows:

1. on the entitlement to *transitions* from perceptual experience to judgements that share that content;
2. entitlement is an evaluative or normative, not merely descriptive concept;
3. entitlement exhibits a truth-connection (hence it is an epistemic notion); and finally,
4. entitlement comes in degrees. (Peacocke, 2004, p. 6-11)

A further reason to think that what Peacocke says about entitlement could be extended to justification (at least as I have been explicating it) is that Peacocke thinks of an entitled transition as a ‘rational’ transition, where this is understood as being incompatible with a *purely* reliabilist account. Peacocke explicitly states that he is “opposed to what is sometimes called ‘externalism’ about justification in epistemology.” (Peacocke, 2004, p. 12) He thinks that the case against pure reliabilism is in part made by examples like cases of clairvoyance that I discussed earlier. (See Chapter 2)

However, other than saying that rationality is not (merely?) a matter of truth-conduciveness, Peacocke gives no positive account of what epistemic rationality consists of. Nevertheless, we can evaluate what he *does* say about entitlement and truth-conduciveness. In addition, however, I suggest that the best explanation of the fact that Peacocke goes on to provide the *a priori* arguments I will discuss below is that he thinks that the availability of such arguments is in part what is necessary for distinctively *rational* truth-conducive transitions.

According to Peacocke, one is entitled to form a certain belief if the belief is the product of a transition one is entitled to make; a partial answer to what comprises an entitled transition is captured by Peacocke's first principle of rationalism:

Principle I: *The Special Truth-Conduciveness Thesis*

A fundamental and irreducible part of what makes a transition one to which a thinker is entitled is that the transition tends to lead to true judgements (or, in case the transition relies on premises, tends to do so when its premises are true) in a distinctive way characteristic of rational transitions. (Peacocke, 2004, p. 11)

This principle in effect states a necessary condition for entitlement, and in so doing gives an account of the truth-connection: truth-conduciveness is a necessary property of entitlements; what else is required, such as what it is for a transition to be rational, he does not say.

Almost immediately after stating this principle, Peacocke modifies it to account for the New Evil Demon intuition I defended in Chapter 3. In such 'demon worlds', Peacocke recognizes that one is entitled to various beliefs, despite the fact that the demon ensures that the transitions one makes are not (actually) truth-conducive. The

requirement that the transition be truth-conducive in a distinctively rational way is to hold not in the actual world, but “only in worlds of a kind which one has a *prima facie*, defeasible entitlement to believe one is in.” (Peacocke, 2004, p. 13) The parallel with Goldman’s Normal Worlds Reliabilism is clear.

However, unlike Goldman’s externalist approach, the cogency of the Special Truth-Conduciveness Thesis will rest on the cogency of the argument for our entitlement to believe we are not the victims of an evil demon. Without such argument then at best we are left with *de facto* conditional reliability, which by Peacocke’s lights, is not sufficient for entitlement. What Peacocke’s Internalism amounts to, I suggest, is maintaining that a necessary condition of entitlement holding is the *a priori* availability of a sound argument that provides a justified presumption in favour of believing that we are in a ‘normal’ world, thus establishing that conditional reliability is likely actual. But before I evaluate the arguments Peacocke provides, what is the *motivation* for conceiving of the truth-connection in terms of *truth-conduciveness*?

Motivations for the Special Truth-Conduciveness Thesis:

The first motivation offered for this principle seems to be a conceptual truth about the nature of judgement. Peacocke writes, “It is a constitutive aim of judgement that one tries to judge that *p* only if it is true that *p*.” (Peacocke, 2004, p. 13) So far, so good -- here Peacocke seems to be endorsing a version of a point that is more commonly expressed by saying that ‘belief aims at truth’. But given this, it is not clear how the next line follows. Peacocke continues, “Rational ways of coming to make judgements must be ones that tend to lead to the truth.” (Ibid.)

First, this second sentence is incompatible with Peacocke's own view that in the demon world one is still entitled to one's beliefs – the way these beliefs were formed does not tend to lead to truth nor affect their rationality. Second, how could a *tendency* for truth follow as a necessary condition from an *aim* of judgement? By analogy, an archer need have no less *aimed* at a target just because he misses his target. In fact, an archer could still aim at the target despite *never* hitting it, be it either because he is a very poor aim, or because an evil demon inevitably intervenes, and knocks his arrow off course. In general, there is no inference from 'I am doing X to get Y', to 'the only way I am justified in acting is in a Y-conducive way'. Perhaps one has evidence that X is Y-conducive, even if it is not. So while it may be true that the aim of judgement is truth, this does not motivate a notion of truth-conduciveness as the correct account of the truth-connection.

To clarify this further, note that by 'tendency' one could either mean the *frequency* with which something happens, or else the *disposition* for something to occur. These are distinct. Simply because something lacks a tendency in the first sense because it is blocked, e.g. rational ways of coming to make judgements might not *actually* lead to truth due to the meddlings of an evil demon, it does not follow that it lacks a tendency in the second sense. For example, a vase may have tendency to break if dropped because it is so *disposed*, even if that disposition is masked in the sense that the circumstances in which the disposition would manifest itself are not allowed to occur, e.g. the vase is wrapped in foam.

Similarly, someone might maintain that rational ways of coming to make judgements must be ones that tend to lead to truth in the sense that such a disposition

exists, whether or not it is masked. It may well be true that dispositions exist, even when masked. However, my point is not that the existence of preventers is necessarily incompatible with the something having a given disposition. Rather, my point is that there is nothing about the *aims* of something that entails any *tendencies*, be they frequencies of something happening, or related dispositions. The aim of judgement may well be *compatible* with rational belief-forming methods tending to lead to truth, but such an aim does not *support* such a conclusion. These considerations *presuppose* the Special Truth-Conduciveness Thesis; they do *motivate* it, which is Peacocke's intention.

Another way of illuminating the issue of whether or not rational ways of coming to make judgements must be ones that tend to lead to truth is to consider the case of induction. One might wonder: does the rationality of inductive inference depend on its being truth-conducive? This is far from obvious, and the traditional problem of induction shows just how difficult it is to convincingly justify that claim that induction is truth-conducive. But that is not to say that induction has *nothing* to do with truth. Even pragmatists tend to hold that induction exhibits a *truth-connection*, which is what in part supports the rationality of inductive inference.

For example, Reichenbach and his followers hold that one cannot prove inductively or deductively that induction will always, or even sometimes, lead to truth.⁷¹ But they think that they can show that *if* there is any method of belief formation that leads to truth, and ultimately knowledge, induction will do so. That is, induction will succeed if any method will, so it is rational to have confidence in induction and its results.

⁷¹ For Reichenbach's original pragmatist response to the problem of induction, see Reichenbach (1938) chapter 5, and Reichenbach (1949) chapter 11.

The argument is very roughly that even though we cannot establish *a priori* the thesis that nature is uniform (which, following Hume, we might assume is a necessary suppressed premise of all inductive arguments), we can nevertheless examine that thesis and discover the following: if nature is uniform and induction is used, it will be successful; if nature is not uniform and induction is used, it will fail. However, if a method other than induction is used and nature is uniform, either success or failure is possible. But importantly, if a method other than induction is used and nature is not uniform, it *must* fail as well. (See Salmon, 1974, p.85-6 for presentation of this argument) This is because if the method did work, it would generate uniformity to which induction itself could be applied. So if that non-inductive method succeeds, then induction succeeds also. In short, “We have, therefore, everything to gain and nothing to lose by using induction. If induction fails, no other method could possibly succeed.” (Salmon, 1974, p. 86)

Of course, pragmatic approaches to the problem of induction are not without their critics.⁷² But what is important here is that induction is often thought of as a paradigm of rational inference (somewhat paradoxically, given the difficulty of justifying induction), and yet there is no assurance that induction is truth-conducive. Yet even by pragmatist lights induction exhibits a truth-connection, albeit a lesser property than the one Peacocke advocates. This reinforces that it is not obvious that rationality or justification must be characterized in terms of truth-conduciveness as Peacocke maintains in his Special Truth-Conduciveness Thesis.

Peacocke goes on to offer three other motivations for his conception of the truth-connection in terms of truth-conduciveness. The first is an argument that a pure

⁷² See for example Lenz (1974).

conceptual role theory of rationality is inadequate to account for the fact that judgement aims at truth. (Peacocke, 2004, p. 15) It is not clear that conceptual role accounts are not sufficient; but even if they are inadequate, it is not obvious that there is not another way of handling the fact that belief aims at truth. Below I present arguments against the truth-connection being characterized in terms of truth-conduciveness, so if one is facing the problem of the aim of belief, one needs another solution than the one Peacocke is advocating here.

The second is a claim that one cannot rule out spurious meanings unless one relies on considerations having to do with reference, semantic value and truth. (Ibid., p. 17)

The third is an argument to the effect that we cannot explain truth and reference in terms of reasons and entitlement. Rather, it must be the other way around. (Ibid., p. 21)

Without developing these arguments further and evaluating them individually, one can easily see that they do not motivate an account of the truth-connection in terms of truth-conduciveness. The considerations Peacocke gives in support of the principle, as he says, implicate “a *role* for a substantial notion of truth and reference in the explication of rationality.” (Ibid.) (emphasis added) That much may well be true. That is, perhaps we must “allude” to truth and reference to elucidate rationality, *pace* pure conceptual role theorists. But alluding to truth and giving it a substantial role in a theory of rationality is not equivalent to that role being one of *truth-conduciveness* (even in its conditional form). So the question remains: what motivates thinking of the truth-connection in terms of truth-conduciveness and reliability, especially given the firm and deeply-rooted intuitions I defended in connection with the New Evil Demon problem in Chapter 3 that count against such a conception? I conclude that there is no such motivation. The only

constraint here is that there *must* be a truth connection if the justification in question is genuinely epistemic. But other than the connection not being one of actual truth-conduciveness, the exact nature of the connection is still an open question.

What if, despite what I have argued above, there *is* some motivation for positing that the connection between justified belief and truth is one of truth-conduciveness? In that case, after providing the motivation, in order to defend the *Special Truth-Conduciveness Thesis* in a way compatible with what I am calling Peacocke's Internalism, one would have to do two things: first, one would have to show that we are entitled to believe that we are in a non-demon world; second, one would have to show that in non-demon worlds the transitions are likely to result in true judgements. This is exactly what Peacocke attempts. I take these attempts as evidence that he thinks that such arguments are necessary for his Special Truth-Conducive Thesis to hold, since by his own admission, reliability on its own is insufficient for entitlement. If he did not think such arguments were necessary for this purpose, it would be a mystery why he modifies the principle in terms of truth-conduciveness only in worlds in which we are entitled to believe that we inhabit, and then go on to offer such arguments.

I will now show that Peacocke's account sets a burden that he is unable to meet, and in fact, the burden may be impossible to meet. This results in its being inconclusive whether or not any of our beliefs are justified. As such, Peacocke's account must be rejected, which in turn further reinforces the conclusions of Chapter 3 that truth-conduciveness (actual, and now, conditional) is not the correct account of the truth-connection.

4.4. Peacocke's *A priori* Arguments Against Scepticism:

Peacocke offers two arguments to the best explanation that aim to refute scepticism and establish the reliability of sense perception, respectively. The arguments are closely related, and in fact share their first two premises. It is perhaps best, therefore, to see the second argument as an elaboration of the first. Peacocke's argument that we are entitled to believe that we are in a non-sceptical world can be summarised as follows:

- 1) Experiences with content are complex. (Peacocke, 2004, p. 86-7)
- 2) A complex phenomenon is more likely to have a complexity-reducing explanation than an explanation that does not reduce complexity or no explanation at all. (Ibid., p. 83; 95)
- 3) Sceptical explanations of our content-bearing experiences do not reduce complexity. (Ibid., p. 90-1)
- 4) Our standard explanations of those experiences do reduce complexity. (Ibid.)

Therefore, we are entitled to believe that we are in a standard, non-sceptical world.

Much in the argument hangs both on the notion of complexity as well as the complexity-reduction principle expressed in premise 2. As Peacocke uses the term, complex phenomena are ones that seem improbable but in fact have an explanation of why they occur, for example, like the structure of a snowflake. (For examples of complexity in his sense, see Peacocke, 2004, p. 75-86) Something is an instance of complexity when the range of exemplified properties is narrow compared to the range of all possible properties. To take Peacocke's example of a snowflake, it exhibits a complex structure, since of all the possible ways it could have been, it exhibits six-fold symmetry, rather than any other shape. Perceptual experience is similarly complex in this sense, since a perceptual experience exemplifies only a narrow range of possible properties that it could have otherwise had.

Premise 2 supposedly embodies a kind of complexity-reduction principle. Peacocke officially formulates the principle as follows:

Complexity Reduction Principle: Other things equal, *good explanations* of complex phenomena explain the more complex phenomena in terms of the less complex; they reduce complexity. (Peacocke, 2004, p. 83) (emphasis added)

This principle, which just states a necessary condition of *good explanations*, is immediately followed by the metaphysical principle that I paraphrased in Premise 2:

Qualified Principle of Sufficient Reason: Other things equal, *it is more probable* that a complex phenomenon has a complexity-reducing explanation than that it has no explanation, or that it has one that does not reduce complexity. (Ibid.) (emphasis added)

Peacocke takes both of these principles to be knowable *a priori*. He tells us that the thought behind these principles is that it is more likely that things come about in easier ways, and that for the most part, it is more rational to believe that things come about in these ways. Peacocke repeatedly emphasizes throughout his discussion that the notion of probability he is interested in is an objective, mind-independent matter.

The third and fourth premises of the argument hold that sceptical arguments are not complexity reducing, and so should be rejected in favour of standard natural selection explanations as the cause of our perceptual experiences. According to Peacocke, a complexity-reducing explanation is one where the phenomenon appealed to has less complexity than the phenomenon to be explained. The problem with sceptical hypotheses is that they allegedly reproduce or multiply the complexity, not reduce it. Typical sceptical scenarios are filled with unexplained explainers, such as issues concerning the intentions of the evil demon, or why and how random events can give rise to our perceptual experiences. So, given the need to explain our perceptual experience, and given the alleged inadequacy of sceptical scenarios, combined with a qualified version of the principle of sufficient reason, Peacocke concludes by inference to the best

(complexity-reducing) explanation that we are entitled to believe that we are in a standard, non-sceptical world.

Before evaluating this argument, I will briefly outline Peacocke's second anti-sceptical argument, since it shares the same problems as his first. The second argument is designed to establish that transitions from perceptual experience to perceptual belief are likely to result in true judgements in non-sceptical worlds. This conclusion aims to elaborate the conclusion of the first argument by showing that not only are we entitled to believe that we are in non-sceptical world, but that the perceptual beliefs we form are usually true. Peacocke's argument for the reliability of sense perception can be summarized as follows:

- 1) Perceptual experiences are complex. (Peacocke, 2004, p. 86-7)
- 2) A complex phenomenon is more likely to have a complexity-reducing explanation than an explanation that does not reduce complexity or no explanation at all. (Ibid., p. 83; 95)
- 3) A natural selection explanation of the occurrence of perceptual experiences is complexity-reducing. (Ibid., p. 87-8; 98)
- 4) It is not clear that any other explanation is complexity-reducing. (Ibid.)
- 5) Therefore, a natural selection explanation of the occurrence of perceptual experiences is probably true. (from 2, 3, and 4)
- 6) A natural selection explanation entails that those experiences are usually veridical.

Therefore, transitions from those experiences to content-endorsing judgements usually result in true judgements. In other words, transitions from perceptual experience to perceptual belief are truth-conducive.

Premises 1 and 2 are the same as Peacocke's first formulation of the anti-sceptical argument. I will grant premises 3 and 4 for the sake of argument. While premise 6 has been questioned, I will not address it here.⁷³ The most problematic premise is one shared by both formulations of the argument: premise 2, the complexity-reducing premise. What does seem both *a priori* and true is the Complexity Reduction Principle that states

⁷³ The main argument that questions if evolution (without God's help) necessarily selects for truth-conducive cognitive faculties is Plantinga's so-called 'Evolutionary Argument Against Naturalism'. For an expression of this argument, and a collection of essays critically evaluating it, see Beilby (2002).

that all else being equal, *good explanations* of complex phenomena explain the more complex phenomena in terms of the less complex. This is a truth about what makes for good explanation. Notice that this is not necessarily the same thing as a *true* explanation. A good explanation may have moral, aesthetic or pragmatic virtues. Peacocke surely recognizes this, which explains his appeal to his Qualified Principle of Sufficient Reason.

Recall that this principle is cast in terms of truth, not in terms of what makes for a good explanation. It states that, all else being equal, *it is more probable* that a complex phenomenon has a complexity-reducing explanation than that it has no explanation, or that it has one that does not reduce complexity. It is this principle that is required for Peacocke's argument to hold. Peacocke claims that this principle is *a priori*. It is not analytic, so presumably it is intended as a synthetic *a priori* truth. Unfortunately, however, it is just not clear if this principle is true, so at best it is inconclusive if the argument is sound.

To see why this is so, consider again the six-step reconstruction of Peacocke's argument above. The overall strategy is as follows: we know, as common sense assures us, that many of our explanations are true. Taking Peacocke's examples, we know that natural selection is the correct explanation of biological evolution, including the evolution and proper functioning of our perceptual faculties. We also take ourselves to know why snowflakes form as they do; to add a further example, we know why certain treatments cure some diseases rather than others. Peacocke then notes that what is common between these true explanations is that they are complexity-reducing in his sense. So from explanations known to be true on one hand, and a disposition to believe that these explanations are complexity reducing on the other, Peacocke reasons that there

is a (justified) *presumption* to believe that complexity reduction connects with truth. What Peacocke is arguing for is a presumption (albeit a defeasible one) that complexity reduction is a sign of truth. Therefore, providing a single counter-example of a correct explanation that *increases* or *reproduces* complexity will not count against such a presumption.

However, the first problem with the Qualified Principle of Sufficient Reason and the use of complexity-reduction is that we cannot know *a priori* the correctness of the explanations we take ourselves to know that Peacocke appeals to. Take the explanations that we know to be true that Peacocke cites as data for the basis of his inference to the best explanation that these (true) explanations are non-accidentally complexity-reducing: natural selection in biology, and the six-fold symmetry of snowflakes. These things are both true, and known to be true (common sense tells us), but they are *NOT* known *a priori*. To be sure, these are paradigms of success in the empirical sciences, not discoveries that could have been made from the armchair. Accordingly, it is difficult to see how Peacocke can gain *a priori* support for his principle on this manifestly *a posteriori* basis.

A second concern is whether we are entitled to take as *known* the data that Peacocke does in this context. Take the perceptual experience referred to in premise 1, for example. Whose perceptual experiences is Peacocke talking about? His own, or the experiences of others as well? If he is including other's experiences too, is he entitled to presuppose the existence of other minds at this stage of the argument? If he is only talking about his own experiences, then maybe the correct explanation of them is not one

in terms of natural selection. How does he know that he has been around to test the value of experiences?

Common sense does assure us that we know that, for example, minds other than our own exist, or that our best evolutionary theory is (at least approximately) true. This is not problematic in an everyday context, or in scientific practice. However, here Peacocke is mounting arguments for the reliability of sense perception, and against the possibility of being a brain in vat – in short, Peacocke is offering abductive arguments against scepticism. Since the falsity of scepticism is what he intends to *prove* we are justified in believing, he cannot *presuppose* the falsity of scepticism by taking it that we know the explanations he appeals to are true, e.g. that there have been millions of years of evolution, without begging the question.

The third problem is that the Qualified Principle of Sufficient Reason is intended in terms of objective, metaphysical probability, but the kind of possibility that Peacocke draws on when trying to garner intuitive support for the principle is epistemic probability. It is certainly true that we *prefer* complexity-reducing explanations, but why is it supposed to be *true* that such explanations are likely to exist? What Peacocke tends to say in support of this are things like the following:

A good theory must not only explain the occurrence of experiences; it should also explain *without extreme improbability* and without pushing the question back why there is the instantiation of just that complex network of relations involved in those experiences having the contents they do. (Peacocke, 2004, p. 97) (emphasis added)

The question is, if the kind of improbability is not epistemic, how do we judge *a priori* if the explanation given is improbable or not? Improbable relative to what? How could we know such a thing? And even if complexity-reduction is ever satisfied, how can we ever

tell that it is *the* simplest explanation (and not just *one* of the simplest explanations)? It is of course easy to tell what *seems* simple and more complex to us, but the question here is justifying the judgement that a given explanation is *in fact* the simplest one.

One possible answer (that would need development and defense) is that conceivability, which is an epistemic notion, is somehow a guide to metaphysical possibility. But even if this is true, is it really *inconceivable* that a complex phenomenon has no explanation, or if it does, it cannot be as complex, (or even more so), than that which it explains? While perhaps unsatisfying, on the face of it there does not seem to be anything inconceivable here, and hence by hypothesis, there is no impossibility. If this is not a reason to think that Peacocke's principle is true, and his stated defense is inadequate, can we find another explanation for why he is misled into thinking that he knows the Qualified Principle of Sufficient Reason *a priori*?

The best explanation of why Peacocke unjustifiably thinks that the Qualified Principle of Sufficient Reason is necessarily true is that he thinks it follows from a related but distinct genuine *a priori* truth. Peacocke asserts: "That it is rational to hold that things have come about in a way in which they are more likely to have come about seems to be an *a priori* principle." (Peacocke, 2004, p. 83) The implication seems to be that it is a rational thing to believe precisely because it is true (perhaps in virtue of meaning alone, no less; or at least knowable *a priori* if the claim is *synthetic*) that it is more likely that things come about in easier, rather than in more improbable or difficult ways. The difficulty is that one ought to be agnostic about what is an easy way for something to come about, given that we are concerned with objective, metaphysical possibilities. This is compounded by the fact that when considering the reliability of sense perception and

the possibility that we may be brains in a vat, we have no idea what the initial conditions are, which renders impossible judgements (let alone *a priori* judgements) about the ease in which things can objectively come about. Further, the point remains that while it is of course easy to tell what *seems* like the simplest explanation, the difficult question here is justifying the judgement that a given explanation is *in fact* the simplest one.

In short, the problem with Peacocke's arguments is that it is simply not convincing that we can know *a priori* that the simple is the sign of the true. But even if it was known *a priori*, difficulties remain in the application of the principle. Without this key premise that complexity reduction is a reason to think something is true, Peacocke cannot *show* that we are entitled to believe that we are in a non-sceptical world. Without this key premise, by Peacocke's own admission, he cannot *show* that the transitions we make are truth-conducive, which in turn undermines his first principle of rationalism in that Peacocke makes our being entitled dependent on being able to *show* that such transitions are likely to result in truth. Without this principle, Peacocke has no explanatory basis for perceptual entitlement, and no account of the truth connection. Worse, since Peacocke's Internalism makes being justified *dependent* on being able to positively show *a priori* that we are in a non-sceptical world, and since he fails to establish this conclusion, he fails to account for the fact that certain beliefs of ours are obviously justified.

In light of the inconclusiveness of Peacocke's arguments, where do we stand? Since Peacocke's Internalism is incompatible with showing that the conditional reliability truth-connection holds, at best we are left with the *de facto* conditional reliability of our belief-forming processes. But given the falsity of externalism, such a relation is

insufficient for justification. But is *de facto* reliability nevertheless necessary for justification? On the one hand, it does not seem to be possible to prove that conditional reliability is *not* present whenever epistemic justification is. On the other, that is not required for our purposes here, which is to explore the plausibility of endorsing the claim that a necessary condition of being justified is the availability, *a priori*, of a sound argument that positively establishes that we are in a 'normal' world (what I have been calling 'Peacocke's Internalism'). But no such argument is available and we are nevertheless justified; therefore, satisfying such a condition cannot be a necessary condition for justification obtaining.

To clarify this, consider a world, CR, where a subject S's beliefs *would* have been reliable, had an evil demon not ensured that they are mostly false. Intuitively, such a world is one where the subject's belief-forming mechanisms are conditionally reliable. Now consider a world, CR*, which is subjectively accessible to S in all the same ways as CR, but we stipulate that S's belief-forming mechanisms are NOT conditionally reliable, i.e. had the demon not ensured that S's beliefs are mostly false, they still would have not been mostly true. Is a world like CR* possible? It does not seem as if it is. Given that the subject is aware of all the same things in both worlds, if a way of forming a belief is conditionally reliable in one world, it must be conditionally reliable in the other.

For example, being conscious of a visual experience as of having hands *prima facie* justifies a belief that one has hands. It is possible that a demon could ensure that forming a belief that one has hands on a basis of that kind of experience might be highly unreliable. But if no demon intervenes, one might surely insist that forming a belief in this way is reliable. Accordingly, forming beliefs in this way is conditionally reliable,

one might maintain. But how could there be a world where experiences of having hands are accessible to the subject, but that forming beliefs about whether or not one has hands on that basis would not likely result in true beliefs, even if the world was normal? Such a world seems impossible.

However, the preceding considerations were raised from a *third-person*, not a *first-person* perspective, which is key to kind of epistemic internalism we are discussing. From the first-person, reflective perspective, we cannot access, prove, or show *a priori* that we are justified in believing that the conditional reliability relation holds, or that we are in non-sceptical worlds. So while things *may* be normal, or conditional reliability *may* be present when we are justified in believing as we do, if we had to have a sound *a priori* argument available to show this as a necessary condition of being justified, then we would lack epistemic justification for most (if not all) of our beliefs. But given that we *are* justified in believing many of the things we do, satisfying such a condition must not be necessary for epistemic justification. This of course presupposes that we are justified in much of what we believe. But this is permissible in this context because I am not arguing against scepticism, but rather, am aiming to elucidate the nature of epistemic justification itself. (For more on presupposing ‘Particularism’ as a methodological constraint throughout this thesis, see the Introduction)

A final set of complications arise that count against accepting conditional reliability as a necessary condition on epistemic justification. Essential to approaches like Goldman’s and Peacocke’s is the notion of a *normal* world, as well as *reliability*. While throughout this chapter I have been charitably proceeding as if these are serviceable notions, it is not clear that either of these notions can be made precise enough

to be of any use. To be sure, Goldman's description of what constitutes a 'normal' world must be false. Recall Goldman's claim:

We have a large set of common beliefs about the actual world: general beliefs about the sorts of objects, events, and changes that occur in it. We have beliefs about the kinds of things that, realistically, can and do happen. Our beliefs on this score generate what I shall call the set of *normal worlds*. These are worlds consistent with our *general* beliefs about the actual world. (original emphasis) (Goldman, 1986, p. 107)

So what fixes the normal worlds is our shared set of common general beliefs about the world. But what if these beliefs were wild, crazy, and unfounded? For example, what if we are considering general beliefs about the world held by primitive people, the uneducated, the insane, and even between different cultural groups? These general beliefs obviously vary widely. Why should we care that a belief would be true in *those* worlds? At best, perhaps beliefs that would be true in those worlds would be blamelessly held, depending on how those beliefs were formed and maintained, but blameless belief is not the same thing as justified belief.⁷⁴

Also, if the beliefs defining normal worlds are quite general, then there will be a whole set of worlds under that description and there will be no guarantee that our belief forming methods will be reliable in all of them. This is because it is just not clear how much should be included in defining a normal world. For example, most people do not have beliefs about whether or not lighting conditions are normal, nor do most people have views about the incidence of hallucination, illusion, etc. So there is no guarantee that in so-called normal worlds our belief forming methods are likely to be reliable. As such, it is hard to see how an appeal to perceived 'normality' will help in assuring truth-

⁷⁴ See 5.3. For a further defense of this claim, see Pryor (2001).

conduciveness (even conditionally). There is a more to ‘normal’ worlds than those that are picked out by beliefs that a given group of people might happen to hold.

Further problems arise if the beliefs that define the normal worlds are too general. If the beliefs are too general, they may pick out an *infinite* number of worlds. The problem with this is two-fold: first, it is unrealistically demanding necessary condition of being justified that the subject be reliable in an infinite number of worlds; second, it would be impossible to judge whether the subject actually is ever justified, since we do not have the ability to judge what would be true of an infinite number of different possibilities. Intuitively, whatever else is true of ‘normal worlds’, there is not an infinite number of them. Also, if the beliefs that define normal worlds are very general, it is likely that at least some of these beliefs will be *inconsistent*. If they are inconsistent, then there will be no world where they are all true. In effect, some of the ‘normal worlds’ as defined will be *impossible* worlds, and as such it undermines the idea that belief is justified only if it is true in ‘normal worlds’, since there may be no such worlds.

Peacocke might be thought to be more convincing on this score since he holds that the requirement that the transition be truth-conducive in a distinctively rational way is to hold not in the actual world, or in a ‘normal’ world as Goldman defines it, but “only in worlds of a kind which one has a *prima facie*, defeasible entitlement to believe one is in.” (Peacocke, 2004, p. 13) The essential difference here is that one must be entitled or *justified* in believing that the world is normal. But other than saying that by definition ‘normal’ worlds are non-sceptical worlds, it is difficult to say anything much more informative than that. As a first attempt, one might think a normal world is one whose existence is mind-independent, which obeys causal laws, and which is filled with material

objects, persons and other animals. But of course different philosophers have denied each of these things. Naturally, the fact that philosophers disagree does not show that no one is justified in believing certain things about the world rather than others, but it does complicate matters by making it sometimes difficult to tell what a ‘normal’ world is, because it is sometimes difficult to tell what one is justified in believing, in part due to the possible defeating effects of one’s awareness of the disagreement of one’s epistemic peers.⁷⁵

But a more serious objection here is that Peacocke’s account makes it nearly impossible to determine which are the relevant worlds in which truth-conduciveness is supposed to hold. Unlike ‘actual’ world reliabilism where it is clear which world reliability is meant to hold in, namely the world the subject occupies, Peacocke makes relevant all the worlds in which one has *prima facie*, defeasible entitlement to believe one is in. But *prima facie* defeasible entitlement is *very* easy to come by. For example, on many accounts of testimony, merely having someone tell me that *p* gives me *some prima facie* justification to believe that *p*⁷⁶, in which case it seems that nearly every world view I have ever been told about becomes relevant on Peacocke’s account. It therefore becomes extremely difficult, if not impossible, to keep track of which worlds’ conditional reliability is supposed to hold in, and so it is difficult to make sound judgements concerning whether or not justification is present.

⁷⁵ Might principles like *Sensitivity* or *Safety* be used as ways of handling conditional reliability and in defining normal worlds? So for example, perhaps a normal world is one in which one’s beliefs are *Safe*? A difficulty here is that both *Safety* and *Sensitivity* are externalist principles, and so it is difficult to see how they could be used to vindicate either *Peacocke’s Internalism* or kinds of internalism that embrace the Awareness Condition. For more on these epistemic principles, see footnote 58 and 83.

⁷⁶ For a detailed defense of a non-reductive account of testimony, see Coady (1992).

So far I have been raising concerns about the question of *where* (in which worlds) one must be reliable if one is to have justified beliefs. But there remains the question of *what* must be reliable.⁷⁷ I have been speaking throughout this chapter of ‘belief-forming processes’ and ‘belief-forming mechanisms’ as what must be conditionally reliable. A *token* method or way of forming beliefs that results in true belief is of a more general *type* of method or way of forming beliefs that can be more or less reliable. However, which types of methods are the relevant ones here? This has become known as the ‘generality problem’ for reliabilism.⁷⁸ A process type can be individuated more or less finely, with the associated reliability varying widely. For example, when I form a belief about what time it is on the basis of what my watch says, which method am I using whose reliability is to be assessed? Is it perception or visual perception or visual perception in good lighting or visual perception in good lighting when sober, etc. etc.? Some of these methods are far more reliable than others, and yet it seems difficult (if not impossible) to non-arbitrarily say which method is the one whose reliability is relevant to whether or not epistemic justification obtains.

If internalism required the availability *a priori* of arguments concerning reliability in normal worlds as part of what is necessary for distinctively *rational* truth-conducive transitions as Peacocke maintains, then the Generality Problem seems to be a problem that needs to be solved. However, since the problem seems either unsolvable or too difficult to address in a non-trivial and informative way, then this counts further in favour

⁷⁷ Scott Sturgeon is a philosopher who distinguishes what he calls the Reliability-Where problem from the Reliability-of-What problem that arise for Process Reliabilism. For a discussion of these problems, see Sturgeon (2000), p. 96.

⁷⁸ See Connee and Feldman (1998) for the development of this objection against Process Reliabilism.

of not employing the notion of reliability as a necessary condition on justification obtaining, on pain of unacceptable scepticism.

In summary, without being able to positively show that we are in non-sceptical worlds, we are at best aware of a *conditional* reason to think that our beliefs are true – a condition that we are not justified in believing is met. And assuming that being able (at least in principle) to demonstrate that truth-conduciveness is necessary for justification, we are left with the unacceptable result that our beliefs only *may* be justified. However, given that at least some of our beliefs obviously *are* justified, we need *actual* reason to think our beliefs are true. But this is not something we can get, given Internalism (of any reasonable form) if truth-conduciveness was indeed necessary for justification.

Therefore, in light of the seemingly insurmountable difficulties associated with conceiving of the truth-connection in terms of truth-conduciveness in either actual or conditional form, I suggest we reject Peacocke's account and look elsewhere for an account of the relation between justified belief and truth. I will propose an alternative account of the Truth-Connection in the next chapter.

CHAPTER FIVE

A POSITIVE ACCOUNT OF THE TRUTH CONNECTION OR WHAT MAKES JUSTIFICATION EPISTEMIC (AND IS JUSTIFICATION DISTINCT FROM KNOWLEDGE)?

5.1. Introduction:

Epistemic justification is different from other kinds of justification. Moral, aesthetic and pragmatic justification (if there are such things), seem to essentially relate to goodness, beauty and usefulness, respectively. What individuates *epistemic* justification, from the other types, is *some* connection to truth. So far we have seen what this connection is not: in Chapter 3 I showed that justification need not be *actually* truth-conducive. In Chapter 4 I argued further that, given epistemic internalism of any reasonable form, the truth-connection cannot even be one of truth-conduciveness in a *conditional* mode. In this chapter I will argue that the connection between justification and truth is *conceptual*.

By a conceptual connection I mean that, in virtue of the concepts involved, there are entailments between the concepts, and that grasping a concept is sufficient to grasp concepts that are conceptually connected to it. For example, there are conceptual connections between the concept *bachelor* and the concepts *unmarried* and *man*. There are entailments between being a bachelor and being an unmarried man --being a bachelor is sufficient for being an unmarried man; likewise, being an unmarried man is necessary for being a bachelor. Furthermore, grasping the concept *bachelor* is sufficient to appreciate these entailments *a priori*.

Similarly, I will argue that in virtue of the concepts themselves, epistemic justification and truth are conceptually connected. That is, certain entailments hold between the concepts of epistemic justification and truth, and that grasping the concept of epistemic justification is sufficient to appreciate these entailments *a priori*.

Specifically, I will argue that epistemic justification is epistemic because it turns on evidence; evidence in turn is epistemic because it is conceptually linked with truth. That is, it is not possible to competently employ the concept of evidence without appreciating that it counts in favour of the truth of that which it is evidence for, just as one cannot competently employ the concept of bachelorhood without appreciating that it applies to all and only unmarried men. Epistemic justification, therefore, is conceptually linked with the truth (via evidence), which is what makes it distinctively epistemic.⁷⁹

Before elucidating these connections, I will motivate and defend the idea that being justified is a kind of epistemic success or accomplishment. I will then proceed to argue that being epistemically justified is a kind of epistemic standing distinct from knowledge, before returning to the conceptual connections between justification and truth.

⁷⁹ Again, my primary focus in this work is the justification of perceptual belief, but in the future I hope to extend this account to other areas of epistemic justification, such as memory and beliefs justified *a priori*.

5.2. On Epistemic Ends/Goals:

It has long been suggested that ‘belief aims at truth’, although there has been much disagreement as to how exactly this metaphorical slogan is to be understood.⁸⁰ One might think that since belief aims at truth, and epistemic justification is a property of beliefs that is also truth-connected, perhaps they ‘aim’ at truth in similar ways. So one might think that the nature of the connection between epistemic justification and truth can be elucidated once the notion of ‘aim’ is understood.

For example, it has recently been suggested that the ‘aim’ of belief is to be understood in *functional* terms, which while still not entirely clear in itself, is perhaps intuitively somewhat clearer in this context, given that talk of ‘function’ is not metaphorical. Intuitively, something’s function is what it is suited or designed for, or even more simply, what it does.

Tyler Burge suggests that by reflecting on the nature of a perceptual system, “it [...] is obvious that it is known *a priori* that the central representational function of a perceptual system is to perceive. This function is *a priori* associated with a representational function (to represent veridically).”⁸¹ (Ibid. p. 508) The “fundamental *representational* good”, according to Burge, is “truth”. (Burge, 2003, p. 506) Because Burge holds that the function of a representational system is to represent veridically, he holds that the function of a system of beliefs is to represent truth, since they are

⁸⁰ Bernard Williams first coined the phrase that ‘belief aims at truth’ (1973). See Velleman (2000) and Wedgewood (2002) for interpretations and defenses of this claim. Owens (2003) argues that belief has no aim.

⁸¹ Alexander Bird is another philosopher who contends that the ‘aim’ of belief is to be understood in terms of function. Although he does not argue for this claim or explain how he knows it, presumably he disagrees with Burge that it can be known *a priori* as he claims that cognitive faculties have essential functions in the way that bodily organs do. See Bird (2007). This is because what function an organ has is surely not knowable *a priori*. I take it that, for example, the heart’s function of pumping blood was an empirical discovery, not something knowable from the armchair.

representational states that in part comprise that system. Burge reinforces this point in the following way, “Understanding what a belief is suffices *a priori* to warrant the view that such a system has a function to represent truths. Beliefs aim at truth, and are defective in a certain way if they are not true.” (Burge, 2003, p. 509)

However, ordinary usage may count against speaking of beliefs as ‘defective’ if they are false. We do talk of measuring instruments, such as thermometers, as defective if they give false readings. But we do not call the readings themselves defective, but merely false or incorrect. Similarly, perhaps perceptual belief-forming mechanisms aim at truth in the sense that it is their function to represent truths, and are defective if they result in false beliefs. The beliefs themselves, while not defective, are false or incorrect. If so, Burge may have the materials to explain the metaphor of the aim of belief in terms of function. I interpret his claim that ‘belief aims at truth’ as meaning that beliefs are individuated by the facts that their *function* is to represent truths, and that the state presents itself as “committal” in a certain way. (Ibid., p. 524) That is, a belief is a mental state whose function is to represent its content as actual, as opposed to merely neutrally or hypothetically.

Much more needs to be said about the aim of belief, but suppose one grants that it is to be understood in terms of function, and that the function of belief is to represent truths, so that beliefs are individuated from other mental states as ones that are essentially truth-directed. Is the function of epistemic justification to be spelled out in terms of truth as well? I will argue that even granting assumptions about the function of belief, there must be a difference in the aim of epistemic justification over belief, since, somewhat paradoxically, there is nothing inherently *epistemic* in a truth aim as such. It is easy to

overestimate how much *epistemic* import arises from belief having a truth-related function, even though truth has *some* essential connection to justification. From the assumption that the function of belief is truth, it does not follow that the function of epistemic justification is truth, even though belief is the state of which justification is a property. As we will see, this is because epistemic justification does not just turn on merely *what* is believed, but rather, *how*, or as we might more naturally say, *why* it is believed.

To see why this is so, consider a suggestion Richard Feldman makes about the essence of an epistemological duty: “Epistemological duties are duties that one must carry out in order to be successful from an intellectual (or epistemological) perspective.” (Feldman, 2002, p. 376) Whether or not there are epistemic *duties*, the same basic point seems true of epistemic norms. Norms govern how something *should* operate. Epistemic norms, therefore, are standards for the proper functioning of the epistemic system. If this is correct, then by identifying epistemic success, as well as what it requires, we will thereby uncover what epistemic norms are, as well as what makes them uniquely epistemic.

On a teleological account, success is measured relative to goals. In the case of epistemology, therefore, epistemic success is a matter of achieving the epistemic goal(s). (Feldman, 2002, p. 377) I intend to show that maximal success consists not of simply achieving the relevant goals, but in achieving them in the right kind of way. Much depends on identifying the correct epistemic goal(s). The standard epistemic goal is the one that William Alston puts as follows, viz., “the aim [of] maximizing truth and minimizing falsity in a large body of beliefs.” (Alston, 1989a, p. 83) The rationale for

this specific goal is not entirely clear, and I also raised a general difficulty for goals that involved maximizing two competing elements in chapter 1.2. If it is true that belief aims at truth, then perhaps there is the goal of believing truly for each belief that one holds. But this is a different goal from one of maximizing truth and minimizing falsehood in a large body of beliefs. If one aimed at believing truly for each belief one held (because belief aims at truth), and happened to form a large number of beliefs that were mostly beliefs, then one would thereby come to achieve the latter goal, but not by aiming at it.

Whatever the rationale for the standard truth-maximizing goal (if there is one), even assuming that the general problem regarding decidability I raised in chapter 1 can be solved, I agree with Feldman that this goal is inappropriate as the epistemic goal in that it gives the counter-intuitive result that one can obtain epistemic success if one stumbles upon truth by dumb luck. (Feldman, 2002, p. 378) For example, one might adopt a true belief because of one's own efforts, as a result of reflecting carefully, the weighing of reasons, or exercising their perceptual faculties with excellence. By contrast, someone might adopt that same belief as a result of guesswork, wishful thinking, stubbornness, bias, etc. The "truth-goal" alone does not have the resources to account for why only the first subject's beliefs are epistemically better than the second.

Likewise, it follows from this goal that a person fails epistemically if she believes falsely, even if she used perfectly reasonable, reliable (albeit fallible) belief-forming methods. (Feldman, 2002, p. 378) As I showed in Chapters 3 and 4, this is false. These considerations show that the epistemic goal needs to reflect that epistemic success is not just a matter of *what* is believed, but *why* it is believed, as I alluded to above. This will ensure that satisfying the epistemic goal is an *accomplishment*, a genuine achievement on

the subject's part, and not just something that can be arrived at by sheer fluke by stumbling upon truth.

A second proposal might then be that the epistemic goal is to have knowledge.⁸² From this goal, therefore, one would attain epistemic success if one has knowledge. This goal is certainly an improvement over the truth goal, since it keeps what is appealing about the truth goal while ruling out its weaknesses. This is so for two reasons: 1) since knowledge is factive, the knowledge goal includes the truth goal, since one cannot know that *p* without *p* being true. So whatever appeal the truth goal has is incorporated in the knowledge goal. 2) The knowledge goal rules out what is undesirable about the truth goal, namely that one could discover it through sheer accident and dumb luck. Knowledge excludes the possibility that the truth of what is believed is entirely accidental.⁸³ Knowledge also meets the desideratum that emerged above, namely, that what is crucial is *why* the proposition is believed, thereby ensuring that satisfying the

⁸² For example, Williamson writes: "Knowledge sets the standard of appropriateness for belief... Mere believing is a kind of botched knowing. In short, belief aims at knowledge (not just truth)." (Williamson, 2000, p. 47)

⁸³ While I take it that it is a platitude that knowledge excludes certain kinds of luck, e.g. it is because of the presence of too much of the relevant kind of luck that one fails to know in a Gettier case, knowledge is compatible with other sorts of luck. For example, I can know that the meeting is at 5pm tonight by reading a posted sign that I just luckily happen to pass. If I had not taken the path I did, I would not have seen the sign, and so would not have known about the meeting. So in a sense, it is pure luck that I believed as I did. Such luck, however, is compatible with knowledge.

While the problem of epistemic luck is a difficult one, the solution will no doubt require relativization to the *method* used in forming the belief in question. For example, one might hold a *Safety* principle which states roughly that a necessary condition on knowing that *p* is that one's belief could not have been easily false while using the same (or very similar) belief forming method.

So one does not know that one sees a barn based on its looking like a barn if the barn is situated in a field of visually indiscriminable fake barns. This is because in *very* nearby worlds in which one believes that one sees a barn because one seems to see a barn, one's belief is false (because one is looking at a fake). Here it is mere luck (of the relevant kind) that one's belief is true, and so it does not constitute knowledge.

By contrast, one knows that the meeting is at 5pm based on reading a sign to that effect since there is no nearby world where one reads a sign about what time the meeting is and that belief is false. It may be luck that one saw the sign one did, but there is *no* luck between forming a belief based on what the sign says and the truth of that belief. In short, intuitively knowledge is compatible with certain kinds of luck, and incompatible with others.

Of course problems remain, including how to non-trivially individuate belief-forming methods. For more on the issue of epistemic luck, see Pritchard (2005). For a defense of the unorthodox thesis that there can be 'lucky' knowledge (i.e. a denial of the anti-luck platitude), see Hetherington (2001).

epistemic goal is an accomplishment. So far, the knowledge goal looks like an appealing way of characterizing epistemic success.

Knowledge is admittedly a kind of epistemic success, but it is not the only kind. To see that it is not, consider cases where everything is going as well as it could as far as the subject is concerned, but what is believed is false. It may turn out that the subject believes that p , is justified in believing that p and bases his belief on those grounds (in a way that is not 'Gettiered'), but that since the world is, so to speak, not cooperating, p is false (say because a sceptical scenario obtains). Here, the subject is doing everything she ought to be doing, and one would know that p if things were as they appear, but since they are not, she fails to know. What goes awry, however, is something beyond the subject's control and powers of perception. It seems that since the only factor that is different in the two cases between where the subject knows and where she does not know has nothing to do with the subject (but rather, just whether or not p is true -- which is entirely up to the world, so to speak), then whether the subject *actually* knows cannot be the only thing that constitutes epistemic success. Being justified is also a kind of epistemic success. What it is to be justified is of course an open question, a question that this thesis attempts in part to answer.

By analogy, consider a practical case of a medical doctor writing a prescription for an ill patient. Under what circumstances is such an action justified? To be sure, whether or not the prescribed drugs *actually* cure the patient is irrelevant to whether the doctor's action was justified. We think that what matters is whether the doctor was reasonable in thinking that her actions would help the patient. Being justified is also a

kind of practical success distinct from the kind of success that would have arisen had the agent's intended outcome been realized.

To further reinforce the point that epistemic success need not consist only in having knowledge, consider a further point Richard Feldman makes. He points out that with regard to some propositions, that sometimes one should suspend judgment on their truth, say for example when one has no evidence for or against them. (Feldman, 2002, p. 378) But by suspending judgment one thereby fails to know. But surely this is the correct thing, the epistemically successful thing, to do. As I noted in Chapter 1, the suspension of judgment is an attitude that is apt for epistemic justification as well. Accordingly, the goal whose achievement defines epistemic success need not be *only* the attainment of knowledge. The attainment of justification is also a state of epistemic success. Again, it is a separate question of what is required to reach that success state.

In light of these kinds of considerations, Feldman characterizes epistemological success in terms of "having reasonable or justified cognitive attitudes." (Feldman, 2002, p. 379) The advantage of this characterization of epistemic success over its rivals is that, unlike the truth-goal, it makes success something the subject can rightly take credit for, that is, the subject is the author of his success. Unlike the knowledge goal, however, satisfying the justification goal is something that can be met entirely by factors the subject is accountable for, unlike the actual truth of the belief, which is up to the world. However, it is not clear how this kind of success is uniquely *epistemological* as it stands. One's belief that, despite being new to golf, one will beat one's much more experienced competitor in the competition, may be the reasonable or justified cognitive attitude to increase the likelihood of winning the game. Such a belief, if justified, is

surely pragmatically justified, although not epistemically justified. As I will argue below, some connection with truth still needs to be drawn to ensure the justification is epistemic; but as we have seen, the connection with truth is not a straightforward one.

Before doing so, however, I will further defend my claim in this section that being epistemically justified is a kind of epistemic success distinct from knowledge. I have been maintaining here, and throughout this thesis, that epistemic justification is distinct from knowledge. However, if recent work by Jonathan Sutton is correct, this is mistaken.⁸⁴ Sutton argues for a surprising and unorthodox thesis, namely: justification *is* knowledge. He claims that there is no concept of epistemic justification distinct from knowledge. I will critically evaluate Sutton's position, ultimately concluding that he fails to establish his bold thesis. I will then return to the question of what makes justification epistemic in the remainder of this chapter.

⁸⁴ See Sutton (2005), as well as his full-length monograph, Sutton (2007).

5.3. Is Justification Nothing More, and Nothing Less, Than Knowledge?

In recent work, Jonathan Sutton argues for a surprising and unorthodox thesis, namely: justification *is* knowledge (Sutton, 2005; 2007). That is, a subject is justified in believing that p iff he knows that p . Sutton further claims that there is no concept of epistemic justification distinct from knowledge. Since knowledge is factive, a consequence of Sutton's view is that there are no false justified beliefs, since only truths can be known. In this section I will summarise and critically evaluate Sutton's arguments for this controversial view. Following Sutton, I will begin by outlining two types of beliefs that do not constitute knowledge but that seem to be justified. I will then survey the concepts of justification he discusses and his identification of two of them with knowledge. At that point I will be in a position to consider Sutton's four arguments for his position that justification is knowledge, concluding that Sutton fails to establish his bold thesis.

Two Kinds of (Allegedly) Justified But Unknown Belief:

1. The Unknown Unknown: Gettier cases and Justified False Belief

The first kind of (allegedly) justified but unknown belief that Sutton considers is what he calls an 'Unknown Unknown' belief: these are those in which the subject is ignorant of his ignorance. Justified False Beliefs and Gettier cases are the paradigm instances of such beliefs. An example of a justified false belief might be one like my belief that there is a dog on the lawn on the basis of its looking to me like there is a dog on the lawn, its sounding to me like there is a dog on the lawn, Stephen's testimony that there's a dog on the lawn, etc., but in fact there is no dog on the lawn, but rather a

cleverly disguised sheep. On other hand, a Gettier case is minimally an instance of *justified*, true, belief, where the subject fails to know what is believed. Gettier cases arise when there is *too much* epistemic luck around.

Here is a version of a standard Gettier case: Al stares out into a field of fake barns, visually indiscriminable from the real thing. In the centre of this field deep in fake barn country is a *real* barn. Suppose Al looks at the real barn amidst the fakes and believes, *that is a barn*. Here Al believes it is a barn, it is true it is a barn, and since he has overwhelming evidence it is a barn (it plainly looks like a barn!), intuitively he is *justified* in believing that it is a barn. However, does Al *know* it is a barn? Surely not. There is too much luck around. His belief could have *very* easily been false. Had he glanced at any of the other numerous barns in the field and formed the same belief, he would have believed falsely. It is just sheer luck (in the relevant sense) that his belief is true.⁸⁵ The standard lesson of such cases is that knowledge is incompatible with such luck. A corollary of such cases is that despite not knowing the proposition in question (e.g. *that is a barn*), surely the subject is *justified* in believing as he does.

A consequence of Sutton's proposal is that Gettier cases are impossible. There will still be cases where a subject fails to know a proposition because there is too much luck around, but these will not be cases of *justified*, true beliefs that fail to be knowledge; hence, there will be no Gettier cases since these just are cases of justified, true, belief that are not knowledge.

Sutton agrees that it is the presence of too much veritic luck that prevents a subject like Al from having knowledge. (Sutton, 2005, p. 361) He goes on to say that the reason we are inclined to call such cases instances of justified belief is because the

⁸⁵ See Pritchard (2005) for a discussion of the kinds of luck (in)compatible with knowledge.

subject *would* have knowledge if it had not been for his bad luck (e.g. if he had not been in fake barn country, he would have known that there was a barn in front of him). Sutton thinks that this shows that justification is parasitic on the concept of knowledge. As he says, “Justification in the relevant sense is perhaps a disjunctive concept – it is knowledge or would-be knowledge.”⁸⁶ (Ibid.) A charitable reading of Sutton here construes him as saying what is called justification in Gettier cases is a disjunctive concept.

2. The Known Unknown

As the label suggests, Known Unknown beliefs are ones where the subject is often in a position to know that he does not know the proposition in question. Nevertheless, it is often taken that such propositions are justifiably believed. Sutton’s two examples of this kind of allegedly justified belief are: i) justified beliefs formed on merely probabilistic grounds, such as “I believe that my ticket won’t win the lottery” when one holds a single ticket in a fair lottery consisting of perhaps a million tickets; and ii) justified beliefs formed as inferences to the best explanation, such as “I believe that evolution is the best theory based on the available evidence.”

Sutton goes on to argue that victims of Gettier cases, those who form beliefs only upon probabilistic grounds, and those who hold beliefs that are inferences to the best explanation that fall short of knowledge, all have beliefs that share the same defect: beliefs formed in those circumstances, or in those ways, are all unjustified. Their beliefs

⁸⁶ Alexander Bird argues that justification is ‘would-be’ knowledge. For details see Bird (2007).

are unjustified precisely because they are not knowledge. How then is Sutton thinking of justification? Sutton considers five candidates.⁸⁷

Five Concepts of Justification:

i) *Warrant*: Warrant is Alvin Plantinga's term of art for whatever makes the difference between true belief and knowledge.⁸⁸ As Sutton points out, this usage assumes that there *is* such a thing as Warrant that can be captured by a reductive analysis, and it assumes that Warrant does not entail truth.

Whatever merits a concept of Warrant might have as an epistemic concept in its own right, I suggest that we look elsewhere for a concept of epistemic *justification*. I suggest that it is perhaps most useful to think of justification as a *candidate* for Warrant, a candidate that is disqualified once we appreciate the force of the Gettier problem. Justification *cannot* be Warrant since justification does not mark the difference between true belief and knowledge – justified, true, belief that fails to amount to knowledge abounds in the form of Gettier cases. As Sutton implies with his paradigms of justified belief that falls short of knowledge, what people are interested in when they are interested in justification, is what is present in Gettier cases. What is present in Gettier cases, however, is not Warrant.

⁸⁷ Why does Sutton consider *these* five conceptions of justification, rather than Reliabilist or other paradigmatically externalist conceptions of epistemic justification? He does not say. In fairness to Sutton, the conceptions of justification he *does* discuss have all been influential and central ways of thinking of epistemic justification in contemporary analytic epistemology (see, for example, Alston 1989a). In any case, I will show that his argument fails.

⁸⁸ See 1.1. for further discussion of Warrant and its relation to other epistemic concepts.

ii) *Justification As Blamelessness*:

Sutton holds that being blameless in believing something is one thing, and being epistemically justified in believing it is another.⁸⁹ He claims that certain epistemic problems have been generated by failing to recognize this fact. For example, Alvin Goldman and others have been worried about cases of cultural isolation or otherwise benighted societies where the subjects use highly unreliable and bizarre belief-forming procedures and thereby form false beliefs, but do so blamelessly (after all, it is not their fault they born when and where they were).⁹⁰ Apparently some have been tempted to claim that such subjects are nevertheless justified.

I think it is correct to hold that justification and blamelessness are not *identical*. For instance, being blameless is surely not *sufficient* for epistemic justification. A subject in a pre-scientific society surely cannot be blamed for believing that he can divine future events by reading sheep entrails; but he is not justified in so believing. However, one might plausibly hold that being blameless is *necessary* for being epistemically justified. That is, if one is to be properly blamed for believing as one does, plausibly such

⁸⁹ We might be thought to be entering murky waters here since it is not entirely clear what blame is. It may be true that we do not have a satisfactory analysis of the concept of blameworthiness, but we can proceed intuitively, since we seem to have no trouble judging instances of things that are to be rightly praised or blamed. That is, we know instances of blameworthiness when we see it.

One salient feature of blame is that it is closely associated with the notion of responsibility. One cannot be blamed for something unless one can be held responsible for it. Responsibility, in turn, seems closely linked to the notion of voluntary control. One cannot be held responsible for an action, for example, if it was not voluntarily performed. If this is right, then a notion of epistemic blameworthiness may well presuppose doxastic voluntarism, and as such, its viability as a concept of epistemic justification may well turn on the extent to which, if at all, our beliefs are subject to voluntary control. For recent discussions of doxastic voluntarism, see for example: Ginet (2001); Feldman (2001); Audi (2001).

The issue of voluntary control seems to be one that differentiates deontological concepts of epistemic justification from justification as blamelessness. For example, one may have a duty or obligation to repay one's mortgage, and one violates that duty by not repaying it. However, this does not entail that one is to be *blamed* for not making one's payments. Perhaps one non-culpably lost one's job due to a financial recession. This example also highlights that evaluations of praise and blame require an element of voluntary control, e.g. one is not to be blamed for not discharging one's financial obligations precisely because one is unable, through no fault of one's own, to meet them.

⁹⁰ For example, see Goldman (1992b).

circumstances *defeat* what one would otherwise be justified in believing. On such an account, being blameless is an enabling condition that allows for the possibility of one's belief being justified. But to be sure, justification and blamelessness are distinct.

A second (pseudo?) problem that Sutton holds is generated by failing to appreciate that blamelessness does not amount to epistemic justification is the so-called New Evil Demon problem (see Chapter 3 for further discussion of the New Evil Demon). Sutton under-describes the problem and wrongly concludes that it only arises given mistaken views about justification and blameworthiness.

As Sutton describes the New Evil Demon problem (NED, hereafter), "On many versions of reliabilism, a thinker globally deceived by an evil demon or who is a brain in a vat will have very few beliefs that are justified since they are not arrived at by the relevant kind of reliable process." (Sutton, 2005, p. 369) It is correct that advocates of the NED case do say that the victims of the demon world are justified in holding many of their beliefs. But this is *not* because they are blameless in believing as they do. They may well *also* be blameless; or, indeed, many of the beliefs they hold may be ones that they can be rightly blamed for holding, e.g. beliefs held in the face of sufficient counterevidence.

The reason the victims of the NED case are justified, it is maintained, is because reliability is not necessary for justification. On this view, if the demon victim's counterpart in the non-demon world is justified in holding a particular belief, then so is his deceived twin. 'Internally' (which I take to mean that of which they are consciously aware; see chapters 2 and 3) identical duplicates are justificationaly identical. What the NED case shows is reliability, or other 'external' factors, are not necessary for

justification. It does not, *pace* Sutton, specifically show anything about the relationship (if any) between blameworthiness and epistemic justification. It merely shows ‘sameness’ of justification between a subject and, for example, her recently envatted counterpart. It is held that both will have as many justified and unjustified beliefs as the other, and each will be as praiseworthy and as blameworthy as the other.

So while there may be interesting links between praise and blame, and justification (as I said above, perhaps blamelessness is *necessary* but not sufficient for justification), I suggest we follow Sutton in setting blamelessness aside as a candidate for justification, while at the same time noting that, contrary to what he says, one can reject justification-as-blamelessness while having said nothing about the New Evil Demon problem since it is a separate issue.

iii) The Deontological Conception:

A deontological conception of justification ties justification to epistemic duties, obligations and permissions. In order to finely delineate different concepts of epistemic justification, the deontological concept of justification is to be kept distinct from the issue of praise/blame. Sutton’s thesis about deontological justification is that it is not distinct from knowledge: a belief is deontologically justified iff it is knowledge. Sutton writes, “I claim that we human beings have a negative epistemic obligation: one ought not believe that *p* unless one knows that *p*, for any proposition *p*.” (Sutton, 2005, p. 364) I will consider Sutton arguments that deontological justification is knowledge below.

iv) *The Evaluative Conception:*

Sutton identifies this sense of justification with knowledge also. A belief is justified in the evaluative sense if it is a “good thing from the epistemic point of view”.⁹¹ Being good from the epistemic point of view, as I have been stressing, is usually thought to have something to do with *truth*. As I have said, the qualifier ‘epistemic’ and some connection to truth is needed to individuate epistemic justification from other sorts of justification, such as pragmatic, moral, or aesthetic. Sutton claims the following about evaluative justification: “I claim that what is justified in the evaluative sense is knowledge – it is knowledge that is the supremely good thing from an epistemic point of view and, unlike perhaps truth, it is epistemic enough to *be* justification.” (Sutton, 2005, p. 365) By ‘being epistemic enough’ Sutton is addressing the issue I discussed above about truth as an end not having much epistemic import in and of itself.

Regarding the issue of internalism/externalism, Sutton notes that his theory is strictly speaking neutral. However, since justification *is* knowledge on Sutton’s view, it is as internal/external as knowledge is. In effect, this amounts to embracing a kind of externalism though since no one (that I know of, anyway) holds a purely internalist theory of *knowledge*. External factors, such as whether or not what is believed is true, or how many fake barns are around for example, will in part determine whether or not a subject knows. So Sutton is committed to the position that external factors are *necessary* for a subject to know, but need not think that such factors are *sufficient*, e.g. he can deny that Bonjour’s clairvoyant knows the President’s whereabouts, but Sutton is committed to holding that the victim of the New Evil Demon fails to have justified beliefs since he fails to have knowledge.

⁹¹ Alston characterizes justification in this way as one its four basic features in Alston (1989a).

Internalism (understood here as the view that ‘external’ factors are neither necessary nor sufficient for justification), whether of a Mentalist sort or as one that embraces the Awareness Requirement, has a chance of being plausible only insofar as justification is distinct from knowledge. An upshot of Sutton’s view is that if he is correct about justification being knowledge, then many traditional forms of internalism are false, including the version I am defending in this thesis.

v) *Justification as Reasonableness: Sutton Concedes Too Much?*

The final concept of justification that Sutton considers is spelt out in terms of reasonableness. Sutton thinks that ordinary usage counts in favour of applying the word “justified” primarily to actions, and applying the word “reasonable” to beliefs whenever and wherever philosophers talk of epistemic justification. We commonly do talk of reasonable but false belief, or the subjects of Gettier cases believing reasonably. So if justification is a matter of reasonableness, it is false that justification is knowledge (since one cannot know falsehoods, and one fails to know in a Gettier case).

Why then, one might ask, is justification not to be understood in terms of reasonableness? Sutton’s argument is shaky here. He claims that justification is primarily meant to be a property of *beliefs*, whereas reasonableness is supposed to apply primarily to *people*, and only derivatively apply to beliefs. Sutton writes,

A belief is reasonable in the circumstances in which it is held if a reasonable person would or could hold it in those circumstances. [...] The notion of a reasonable *person* is understood in terms of knowledge; he is one whose belief-forming faculties and habits (e.g. inferential habits) are such as to deliver knowledge when conditions are right. (Sutton, 2005, p. 373)

Given this account, people can have reasonable but unjustified beliefs when they lack knowledge (say, when they are victims of a Gettier case, or when they believe something false on the basis of good evidence), and reasonableness is not equated with blamelessness (remember the cases of cultural isolation). Sutton intends these concessions to make his view more palatable. However, even reasonable people can have some unreasonable beliefs, unless they are perfectly reasonable, so it is false that a belief is reasonable if a reasonable person would hold it, unless we are talking about perfectly reasonable people.

But in making these concessions, one cannot help but wonder if Sutton concedes too much. Has he in effect conceded that there is a concept of epistemic justification distinct from knowledge? He is right that there is a property that is distinct from mere blamelessness that people's beliefs can enjoy even when those beliefs are false. This property, even if not *identical* with reasonableness, is certainly one enjoyed when the subject is being reasonable. Why is not *this* property, which is not identical with knowledge, *justification*? Sutton's primary reason for denying this is that justification is "supposed to" be primarily a property of beliefs, and reasonableness, while it surely applies to beliefs, is supposed to primarily be a property of persons. But then again, it seems that 'reasonable people' could be defined as people whose beliefs are by and large reasonable, and 'perfectly reasonable people' are those whose beliefs are reasonable without exception, so reasonableness applies primarily to *beliefs* and derivatively to persons.

But what if some philosophers have simply misunderstood the sense in which justification attaches to beliefs? That is, it would surely be objectionable if a theory

allowed that justification were not a property of beliefs in *any* sense⁹², but is it really that objectionable if it applies to beliefs in a derivative rather than primary sense?

In a broadly Moorean spirit, what are we more sure of: i) that justification applies to beliefs in a primary, rather than derivative way; or ii) that there are no justified false beliefs -- that there is no concept of epistemic justification distinct from knowledge? The truth of i) seems far less obvious than the falsity of ii). Therefore if the truth of ii) is implied by the truth of i), the first claim should be rejected, or at least not presupposed as Sutton does.

On the model of justification I am considering, a belief is justified if in the circumstances in which it is held a reasonable person would or could hold it in those circumstances. I take it that it is an open question how the notion of a reasonable person is understood. Perhaps Sutton is right: a reasonable person is to be understood in terms of knowledge; he is one whose belief-forming faculties and habits are such as to deliver knowledge when conditions are right. But even if this is right and justification must be understood in terms of knowledge, it does not follow that justification simply *is* knowledge, which is his primary thesis. Sutton is arguing for that bold thesis, not the weaker one that Timothy Williamson defends that it is 'knowledge first', the primitive in terms of which all other epistemic notions are understood.⁹³

On the other hand, perhaps the notion of a reasonable person is to be understood independently of knowledge, e.g. a reasonable person is one who proportions her beliefs to the evidence (whether that delivers knowledge or not). In that case, reasonableness would be an independent epistemic goal from knowledge and so be a distinct form of

⁹² See 1.1. where I note that justification being a property belief is the first of its four basic features.

⁹³ For a defense of this thesis, see Williamson (2000).

evaluative justification. And if there is an obligation to believe reasonably distinct from the obligation to believe only if one knows (e.g. one might have the obligation to believe only if one's belief fits the available evidence), reasonableness would be a distinct form of deontological justification. In making the concessions he does in acknowledging a notion of epistemic reasonableness distinct from knowledge (in at least the sense of reasonableness not being identical with knowledge, even if it is not more conceptually fundamental than it), Sutton casts doubt on the thesis that epistemic justification is materially equivalent to knowledge before even giving his arguments for it.

With a discussion of epistemic justification now in place, we can turn to Sutton's arguments that identify it with knowledge.

The Arguments:

i) The Assertion Argument:

Following Williamson, Sutton accepts the Knowledge Rule of Assertion: one must assert that p only if one knows that p . Sutton argues that the connection between warranted assertability and knowledge is mysterious unless there is the same connection between justification and knowledge. That is, "one must: believe p only if one knows p , [is] a norm that is as constitutive of belief as Williamson's knowledge rule is for assertion." (Sutton, 2005, p. 374)

However, if sound, this first argument shows that the standards governing assertion need to be the same as those governing belief. This in itself is not surprising if the rules governing assertion are simply the 'external' counterparts of rules governing belief. It is a further question what norm(s) do govern assertion. In other words, the

cogency of this argument depends in turn on the cogency of arguments for the knowledge rule.

Sutton embraces Williamson's three arguments for the knowledge rule of assertion. Williamson's arguments are as follows:

- 1) The knowledge rule is allegedly needed to explain what goes wrong in (quasi-) Moorean paradox. The hypothesis is that the reason it is improper to assert " p and I don't know that p " is because by asserting p one implies that one knows it, which is then denied in the second conjunct, which is contradictory. (See Williamson, 2000, p. 253-5)
- 2) It is claimed that the knowledge rule is needed to explain the unacceptability of assertions that one will lose in Lottery Paradox cases. It is allegedly improper to assert that one's ticket will lose the lottery, no matter how likely that fact is. The best explanation of the source of this impropriety, it is suggested, is that one does not *know* that one will lose the lottery solely on the basis of its being highly improbable. (See Williamson, 2000, p. 249-252)
- 3) The final argument for the knowledge rule is, as Sutton says, "simply the appropriateness of challenges and rebukes to assertions such as 'How do you know?' and 'You don't know that!'" The fact that appropriate challenges to assertions involve questioning the asserter's knowledge provides some evidence of the correctness of the Knowledge Rule. (See Williamson, 2000, p. 249-252)

As I said, this argument for the thesis that justification is knowledge presupposes the correctness of the Knowledge Rule of Assertion. If some other principle better explains

the phenomena of assertion than the knowledge rule, this argument fails to go through. I will consider such a principle below.

ii) The Lottery Argument:

The Lottery Paradox can be formulated in terms of knowledge or justified belief. The idea is that if one can know or be justified in believing that that one's ticket will lose the lottery, then, given other intuitive assumptions, one is led to paradox. Following Dana Nelkin, Sutton formulates the paradox(es) as follows:

The Knowledge Paradox

1. Jim knows that his ticket t1 will lose.
 2. If Jim knows that his ticket t1 will lose, then he knows that t2 will lose, he knows that t3 will lose...and he knows that t1,000,000 will lose.
- So,
3. Jim knows that t1 will lose...and Jim knows that t1,000,000 will lose. (1,2)
 4. Jim knows that either t1 will not lose...or t1,000,000 will not lose.
 5. Propositions of the following form comprise an inconsistent set: (a) p1...(n) pn, (n+1) not p1 or...not pn.
- So,
6. Jim knows propositions that form an inconsistent set. (3,4,5)
 7. It is not possible to know propositions that form an inconsistent set.
- So,
8. (1), (2), (4), (5), or (7) is false. (Sutton, 2005, p. 378)

The Justification Paradox shares all the above premises, except that they are cast in terms of 'justification to believe' rather than knowledge. (See Sutton, 2005, p. 378-9)

Again following Nelkin, Sutton argues that the first premise of the paradoxes is the one that must be rejected, i.e. one does not know or justifiably believe that one's ticket is a loser. Sutton departs from Nelkin in his explanation of why the first premise of the justification paradox should be rejected. He maintains that the best explanation of why the first premise of the justification version of the paradox needs to be rejected is because justification is knowledge; so once we see that the knowledge premise must be

rejected in one version of the paradox, we must reject what is in fact the identical, though disguised, premise in the second paradox.

A presupposition that Sutton shares with Nelkin is that we need one unified solution to apply to both paradoxes. Simply put, they in effect argue that since premise 1 should be denied in the knowledge case, i.e. it is allegedly obvious that one does not know that one's ticket will lose the lottery, the first premise of the justification version of the paradox must be rejected also, i.e. contrary to one's intuitive judgement to the contrary, one's belief that one will lose the lottery is *unjustified*. If one rejects this intuitive premise, then one must explain why our intuitions lead us astray. But one might wonder, why accept the presupposition that the Lottery Paradoxes need a unified solution?

Sutton claims that, "it is a clear desideratum of a solution to one version [of the lottery paradox] that it can be applied to the other." (Sutton, 2005, p. 379) It is not clear why there would be such a desideratum, especially since as Sutton admits, the knowledge version of the lottery paradox is not even a paradox! As Sutton says,

Denying premise 1 [that one knows that one will lose the lottery] seems to be by far the best strategy with respect to the knowledge paradox, and Nelkin advocates denying 1* [that one is justified in believing that one will lose the lottery] on that basis. Indeed, it is not remotely unintuitive to deny 1; the knowledge version of the paradox is not really a paradox at all, but simply an unsound argument. (Sutton, 2005, p. 380)

Given this, why think that a non-paradox and a paradox deserve the same treatment independently of thinking that justification = knowledge? As Sutton says, "The identification of justification and knowledge receives support from the fact that it enables a parallel diagnosis of the knowledge paradox and the justification paradox." (Ibid.) But once we see that there is no reason to accept such a presupposition, and there is reason

against accepting it, i.e. it entails that one must deny the common sense view that one is justified in believing that one's ticket will lose the lottery, Sutton's argument that justification = knowledge is undercut.

iii) *The Modesty Argument:*

The third argument Sutton gives for the thesis that justification is knowledge is what he calls the Modesty Argument. If sound, this argument is supposed to establish that there are not any *known unknown justified beliefs*, e.g. justified beliefs formed on merely probabilistic grounds or justified beliefs that are inferences to the best explanation.

Sutton argues that we are often overconfident by categorically believing propositions that our evidence merely supports as probable. For example, even if one cannot know or be justified in believing that one's ticket will lose the lottery, one *can* know or justifiably believe that one's ticket will *probably* lose. Sutton writes, "One gains nothing – no motive for rational action, and no benefits thereof – from believing immodestly that one will lose the lottery that one does not gain from believing modestly that one will very likely lose the lottery." (Sutton, 2005, p. 383) But since we cannot and should not be maximally modest in our beliefs, what marks the difference between an improper and proper level of commitment to a given proposition? Sutton's answer, of course, is that *knowledge* makes the difference between rational and irrational belief formation. (Ibid.)

We can all agree that our beliefs should be as modest as necessary without being under confident in what we believe. It is just unclear why the relevant normative

standard here is knowledge. Take an example Sutton offers that is supposed to illustrate why we should not be maximally modest in our beliefs. “Never believe that one has lost the lottery, simply that it is very likely that one did, even when one watched the draw take place and knows that one lost.” (Sutton, 2005, p. 383) Sutton makes similar remarks about believing a theory that is supported by an inference to the best explanation and beliefs based on perceptual experience.

If these examples show anything, they seem to show that beliefs formed on merely probabilistic grounds or by inferences to the best explanation are insufficient (and in some case unnecessary) for positive epistemic status to obtain. In Sutton’s example, one’s belief that one will lose the lottery is *not* based on its low probability of winning, but rather on one’s witnessing the draw and *seeing that* one lost, indicating that one’s former grounds were insufficient for the belief’s being properly held. But is it really *knowledge* here that governs whether this belief is rationally formed?

I suggest not. Compare two cases. First, suppose one witnesses the draw and hears the announcer call out someone else’s number. It is true that one’s ticket has lost; in the absence of defeaters, plausibly one knows on this basis that one’s ticket did not win. Suppose again that one witnesses the draw and hears the announcer read out a ticket number other than one’s own as the winner. As it happens, the dyslexic announcer read one’s number backwards; one has won the lottery but does not know it yet. Before the result is sorted out, does one justifiably believe that one has lost the lottery on the basis of the announcement? Surely one does. One has exactly the same evidence as one does in the first case. But one fails to know in the latter case since what one believes is false. Since one can justifiably believe without knowing, Sutton’s thesis is false.

What Sutton highlights in the Modesty Argument is not that knowledge governs rationally confident belief formation, but a kind of scepticism for certain kinds of beliefs based on certain kinds of grounds, as well as the fact that we need to be careful not to be overconfident in what we believe given our grounds. Substituting ‘justification’ for ‘knowledge’ in this section gets the same results for rational belief formation even if we stipulate that what is believed is false, showing that we can have justification without knowledge, which falsifies Sutton’s thesis.

iv) The Posterior Evaluation Argument:

Take this claim from Sutton:

If a belief that p is one that would be justified were one to form it, and it is in one’s interest to have a belief in whether or not p , and one is capable of forming such a belief, then, in some intuitive sense, one should believe that p , although that ‘should’ is not purely a ‘should’ of epistemic obligation. It is rather a ‘should’ generated by prudential considerations of self-interest interacting with the epistemic goals that determine an evaluative notion of justification. (Sutton, 2005, p. 384)

Sutton claims that if we substitute ‘would be justified’ with ‘would constitute knowledge’ then the above claim is clearly true. His thesis is that this claim is not true if any other epistemic concept other than knowledge is inserted in its place. I will show that this thesis is false.

This thesis is argued for by means of a variation on the original Gettier case. In Sutton’s case one has the desire to drive a Ford and will receive an incentive for doing so. One has lots of evidence that one’s colleague owns a Ford, but this evidence is misleading. However, one does have another colleague who owns a Ford, so if one formed the belief that a colleague owns a Ford, one would have a justified true belief that

would fail to be knowledge. However, one fails to take notice of one's evidence, and so fails to form any beliefs about who might own what car.

Suppose also, Sutton continues, that one later comes to appreciate the implications of one's earlier evidence, but also learns that the evidence was misleading, so that if one had formed a belief on its basis, one would have formed a true belief that would have been 'Gettiered', so to speak. Sutton now asks, "Is there an intuitive sense in which one will *now* judge that one should *then* have formed the belief that one had a colleague who owned a Ford – an allegedly justified true belief?" (Sutton, 2007, p. 58) Sutton says it is clear that there is, and he is right about one particular sense in which one should have formed the belief. One should have then formed the belief because one would have gotten the incentive. And it is true that that same non-epistemic sense in which one should have formed the belief is present in a version of the case where one forms the belief about the colleague's Ford in the absence of any evidence whatsoever. A true belief not based on any evidence at all would have equally gotten the incentive. However, only in the misleading evidence case, and not the absence of evidence case, is there an *additional* sense to the one that Sutton acknowledges in which one should have formed the belief, viz., that was the well-evidenced belief to form.

Sutton thinks that there is a 'partially epistemic' sense of a 'should' that arises only in a case where the original evidence was *not* misleading and so the subject would have gained knowledge. In such a case, Sutton comments as follows:

Evaluating one's former cognitive situation after the fact, is there an intuitive sense in which one should have formed the belief apart from the consideration that it would have been true? I contend there is: one should have known that one has a colleague who owned a Ford, and this is to say, that in the case at hand, that one let oneself down doxastically speaking in a way that

goes beyond having merely failed to form a true belief. (Sutton, 2007, p. 59)

I agree that in this case there is a sense that one should have formed the true belief, but it is the *same* partially epistemic sense present in the misleading evidence case. One failed to appreciate one's evidence and proportion one's belief to it. In the misleading evidence case, one should have believed that one's colleague has a Ford, and one lets oneself down doxastically by failing to proportion one's belief to the evidence. So contrary to what Sutton says, there is an intuitive sense in which one should have formed the belief that one's colleague has a Ford in the misleading case apart from the non-epistemic sense in which believing truly would have been beneficial. The intermediate case between knowledge and the merely pragmatic case is the one where one should have formed the belief because that is what the evidence supported (even though the evidence was misleading). The existence of the intermediate case shows that justification is not knowledge.

On the face of it, these considerations also count against knowledge norms of belief and hence the Knowledge Rule of Assertion, since this is a case where one should believe that *p* even though one does not know that *p* -- in which case it is false that one should believe (assert) only what one knows.

The Main Objection:

Towards the end of his defense, Sutton discusses an objection, that if sound, devastates the first three of his arguments, and given that I have provided sufficient materials above to undercut his Posterior Evaluation argument, the objection becomes even more pressing. The objection is that a weaker norm-governing assertion will do all

the work that the Knowledge Rule is supposed to. That is, if a weaker norm can handle the quasi-Moorean Paradox cases, the Lottery Cases, and the cases that involve challenging or rebuking assertions, then the Knowledge Rule is unmotivated. This would be devastating for Sutton as the Knowledge Rule is an explicit premise of his assertion argument, and his second and third arguments depend on it as well, albeit less directly.

Sutton considers that the best and simplest candidate to possibly replace the Knowledge Rule is:

(The J-Rule): One must: assert p only if one has a justified belief that p . (Sutton, 2005, p.388)

Sutton thinks that the J-Rule would handle the relevant phenomena that the K-Rule is supposed to, plus (and I think this is a virtue of the J-Rule) it implies that certain false assertions are warranted. However, Sutton rejects the J-Rule because he thinks it falls prey to an anti-complexity argument, as I will explain below. However, despite what Sutton claims, it is not clear how the J-Rule accounts for the data that it is supposed to.

One might think that a J-Rule handles all the cases it is meant to in the following way: first, if someone asserts ' p and I don't know that p ', one might say that that is inappropriate since by asserting p the subject implies he is justified in believing it, but then implicitly denies this. However, if justification is distinct from knowledge, then one can have justified beliefs that do not amount to knowledge. So to admit one's ignorance is not to say one is not justified. There is no contradiction.

So although the J-Rule makes sense of the incongruence of asserting ' p and I am not justified in believing that p ', it does not account for the tension in asserting ' p and I do not

know that p '. So it seems that the stronger Knowledge Rule has an advantage over the Justification Rule.

But it is important to distinguish two kinds of Moorean data:

- 1) Assert: ' p but I don't know that p ' (because the subject has no justification to believe that p)
- 2) Assert: ' p but I don't know that p ' (even though the subject has a highly justified belief that p)

It is the latter datum that the advocate of the J-Rule must explain, since this is where the J-Rule and the K-Rule come apart in what they are able to account for. The J-Rule can only accommodate the former datum.

A similar point holds in the challenges cases, i.e. 'How do you know?' and 'You don't know that!' In making the challenge, e.g., in a lottery case, the challenger may recognize that there is a huge amount of justification for S to believe that p but still feel able to properly challenge the assertion that p by questioning whether they know that which they are asserting. Lottery cases are supposed to be paradigm cases whereby a subject's belief that his ticket will lose is extremely probable and yet he should not assert that his ticket will lose. Here we have justification without knowledge, and indeed a case where one cannot explain the data with the J-Rule since the J-Rule predicts that one can assert that one will lose the lottery. But intuitively this is wrong, since such assertions, while justified, are inappropriate nevertheless.

In light of this, it seems that the best and simplest candidate to replace the Knowledge Rule and the J-Rule is one that links justification and knowledge:

(JBK-Rule): One must: assert *p* only if one has justification to believe that one knows that *p*.⁹⁴

Such a principle explains what is wrong in with the quasi-Moorean and Lottery cases, as well as why assertions can be challenged by challenging the subject's knowledge. In each case, the subject lacks justification to believe that he knows what he is asserting. Unlike the J-Rule, the JBK-Rule can handle *all* the cases that the K-Rule can; but unlike the K-Rule, it allows that subjects are sometimes warranted in asserting what they do not know, e.g. in Gettier cases, as well as cases of justified but false belief. The JBK-Rule, therefore, shares all the explanatory virtues of both the K-Rule and the J-Rule, without sharing their vices.

Why then does Sutton accept the Knowledge Rule over rival rules of assertion, such as the JBK-Rule? It is because he accepts a version of Williamson's complexity argument that he thinks that the K-Rule has a presumption in its favour over other rules since it provides a simpler account of warranted assertion.

The J-Rule is allegedly more complex than the K-Rule because justification is allegedly parasitic on knowledge. Sutton would therefore think that the JBK-Rule is more complex than the K-Rule, since it relies on both the concepts of epistemic justification and knowledge. Recall from the beginning of this discussion where I noted that Sutton says that the reason we are inclined to call Gettier cases instances of *justified* belief is because the subject *would* have knowledge if it had not been for his bad luck (e.g. if the subject had not been in fake barn country, he would have known that he was looking at a barn). As Sutton says, "Justification in the relevant sense is perhaps a

⁹⁴ Williamson himself considers and rejects a similar principle, except casts it in terms of 'rational' rather than justified belief, as well as in terms of doxastic justification (rationality) rather than in terms of propositional justification (rationality) as I do here. See Williamson (2000) p. 260-63 a discussion of this similar principle as well as other belief-based accounts of assertion.

disjunctive concept – it is knowledge or would-be knowledge.” I suggested above that perhaps justification can be spelled out independently of knowledge, e.g. in terms essentially tied to evidence and not to knowledge. If such an account can be given, then the J-Rule is just as simple as the K-Rule. However, now we see that the best rule of assertion, the JBK-Rule, does invoke the notion of justification, as well as knowledge. Perhaps also justification is parasitic on knowledge. What then follows the thesis that justification *just is* knowledge?

If knowledge is the primitive in terms of which all other epistemic concepts must be understood, does that count decisively in favour of the K-Rule? Of course not. A complexity argument is not a definitive argument for anything. Rather, simplicity provides a (defeasible) *presumption* in favour of one theory over the other. A theory needs to have more virtues than mere simplicity, and sometimes therefore, the more complex theory is to be preferred over the simpler one. At least one other essential thing a philosophical principle must do is deal adequately with the relevant cases.

As Williamson himself notes, “Much of the evidence for the knowledge account [or any account of assertion, for that matter] comes from the ordinary practice of assertion.” (Williamson, 2000, p. 243) It is for this reason that Williamson endorses the K-Rule over the admittedly more simple truth rule of assertion (one must assert *p* only if *p* is true). (Cf. Williamson, 2000, p. 242; 262) The Truth Rule does not explain what is wrong with the Lottery cases, or cases where a subject asserts that *p* truly but has absolutely no evidence that *p*. So the more complex K-Rule is motivated via these considerations. However, our ordinary practice of assertion includes not only the Moorean and Lottery cases, but also Gettier cases and cases of justified false belief. If

possible, therefore, we need a rule of assertion that gives a unified treatment of all these cases.

It is here that the JBK-Rule gets support over the K-Rule. The JBK-Rule handles all the cases the K-Rule does, *plus*, it allows for warranted false assertions, something very intuitive that the K-Rule prohibits. In short, since the JBK-Rule does what the K-Rule is supposed to do, and more, by handling a wider range of cases, it should be preferred over the K-Rule, even if it turns out to be true that knowledge is more conceptually fundamental than justification.

Without the K-Rule, by Sutton's own admission, three out of four of his arguments for the thesis that justification is knowledge fail. By Sutton's own admission, the best (and single?) reason to accept the K-Rule over weaker principle-governing norms of assertion is that it is simpler than its rivals. It is not obvious that it is simpler than principles that invoke notions of justification that are not parasitic on knowledge, e.g. justification supervenes on evidence, or justification is a matter of reasonableness (understood independently of knowledge). But even if the K-Rule is simpler than its rivals, there is a principled reason to accept the JBK-Rule over the K-Rule, viz., that it is intuitive that there is nothing wrong with a false assertion over and above the false belief that it expresses. We judge that we can be warranted in making false assertions. In short, the JBK-Rule of assertion explains the phenomena that motivate the K-Rule, while better handling a wider range of cases.

Justification is not Warrant in Plantinga's sense, nor is it mere blamelessness, but there do seem to be deontological and evaluative forms of justification that seem to

amount to a concept of epistemic justification that might be elucidated in terms of reasonableness and/or evidence.

As for the Posterior Evaluation argument, I hope to have shown that there is an intermediate case between knowledge and the merely pragmatic case. That is one where one should have formed a belief because that is what the evidence supported (even though the evidence was misleading). The existence of the intermediate case, coupled with the arguments above in support of the JBK-Rule over the K-Rule, shows that justification is not knowledge.

5.4. The Conceptual Connections Between Justified Belief and Truth:

As I have argued, one of the ways of achieving epistemic success is by having epistemically justified beliefs, and this is kind of epistemic success distinct from those beliefs amounting to knowledge. A necessary condition of being epistemically justified, I argued in Chapter 2, is that the Awareness Requirement is met. It holds that:

S is justified in believing that p only if

- i) there is something, X , that contributes to the justification of S's believing that p ; and
- ii) For all X that contributes, S is aware (or potentially aware) of X .

In Chapter 2 I defended the second conjunct of the Awareness Requirement. But the first conjunct is no less important – there must actually be something contributing to the subject's justification. When considering contingent empirical propositions about the external world, merely being in a conscious state that includes the content that p is insufficient to justify the belief that p . Hoping, wishing and visually imagining that p have the content that p , but they do not justify the belief that p . But we think that a perceptual experience that p *does* justify the belief that p . However, this cannot be purely in virtue of its content since p is not a good reason to think that p is true -- this is the paradigm case of objectionable circular reasoning. Furthermore, such an approach would make perceptual justification conclusive. If one's reason to believe that p were just a perceptual experience with the content p , one's reason would be conclusive since p *entails* p . But this misconstrues the nature of perceptual experience -- experience need not be *conclusive* evidence for belief, because having a (non-factive) perceptual experience does not entail the truth of what is represented, and yet such states nevertheless yield justification, as careful consideration of cases reveals. For example, a brain in a vat is justified in believing that he has hands on the basis of his perceptual

experience as of hands. (See Chapter 3 for an extended defense of this claim) Therefore, over and above its mere content, there is something special about perceptual experience that makes it justificational relevant in a way not shared by other propositional states.

What one is aware of when one has a justified belief, among other things, is *evidence*. That is, the norms that govern how a well functioning epistemic system operates dictate that the subject should have well-evidenced belief that satisfies the Awareness Condition. My contention is that epistemic justification is epistemic because in such cases the subject has reason to think her belief is *true*. Well-evidenced belief is similarly epistemic because in such cases the subject has reason to think that what she believes is *true*. The question that arises then is the following: why is evidence reason to think that what is believed on its basis is true?

The answer to this question is entailed by conceptual truths captured by the following principles:

Truth-Connection: Necessarily, for a belief to be epistemically justified by a consideration, that consideration must count (to some degree) in favour of that belief being true.

Evidence: Necessarily, evidence is that which (*prima facie*) justifies belief.

If only things that count in favour of truth justify belief, and evidence is that which justifies belief, then it follows that evidence counts in favour of the truth of what is believed. Evidence *must* be truth-connected since we confidently judge that belief can be justified on its basis; epistemic justification is epistemic because it is truth-connected in virtue of being based on evidence.

The truth of these principles can be shown by reflecting on our epistemic practice. For example, imagine that I believe that am supposed to attend a party tonight at 7:00, and base my belief respecting the starting time of the party on the information printed on

the invitation. Surely my belief is justified. Suppose further that the host of the party rings me and tells me that the party will actually start at 9:00, and I believe him. Again, obviously this belief is justified. What justifies me in my first belief is my reading the invitation. What justifies me in my second belief about the party starting time is a piece of testimony I received from the host. In each case, visual perception and testimony, the sources justify me in believing as I do because these are considerations that count in favour of the *truth* of the proposition in question. It does not matter if what I read is useful to me, or convenient, or pleasing – only truth-related considerations make an epistemic difference. We would judge that I was epistemically *unjustified* if I maintained my belief that the party was at 7:00 after receiving the host's call on the grounds that believing the party was at 9:00 would upset me as it conflicted with a previous engagement. The best explanation of these judgements is that only considerations that count in favour of truth can make an epistemic difference.

In support of the principle that evidence is that which justifies belief, consider again our epistemic practice. The fact that it is incoherent to claim to be interested in whether or not a belief is (epistemically) justified, but not to be concerned with what (if any) evidence supports it, supports the claim that evidence is that which justifies belief. Considering the example above of the party's starting time, it is telling that there can be no change in what one is justified in believing without a change in evidence (again, see chapters 2, 3 and 4). This suggests it is evidence that fixes what one is justified in believing. Further support for this can be found by noting *what* a defeater defeats. In the case of my belief that the party starts at 7:00, the justification for this belief is defeated by learning something that rebuts what I read on my invitation. A justified belief is defeated

by undercutting or rebutting the *evidence* that it is based on, which implies that it is the *evidence* that justifies belief.

In short, if only things that count in favour of truth justify belief, and evidence is that which justifies belief, then it follows that evidence counts in favour of the truth of what is believed, which gives us the resources to explain what makes epistemic justification distinctly epistemic. Consider cases of perceptual belief generally. In such cases, the subject is justified by basing her belief on *evidence* of which she is aware (conscious perceptual experience). Being aware of such evidence is reason to think that what she believes is true, which is what constitutes the truth-connection, and thereby makes the practice epistemic. Evidence is reason to think that what is believed is true because belief can be justified on that basis, and only truth-connected things can justify belief.

So what is it about perceptual experience, rather than other states, that counts towards a belief sharing the same content being true? The difference is this. Given that for a belief to be justified by a consideration, it must count (to some degree) in favour of that belief being true, and given that evidence is that which justifies belief, it is clear that perceptual experience is evidence for perceptual belief. As such, it counts towards a belief's being true, whereas hoping, wishing and visual imagining do not, as these latter states are not evidence, which is obvious from the fact that belief cannot be justified on their basis. If we had reasons to think that these latter states *did* count towards what they present as being true, they would be evidence, and as such, justificationaly relevant.

Also, perceptual experience necessarily represents its content assertively, rather than neutrally or hypothetically as other propositional states do. A state that presents

itself neutrally or hypothetically cannot be a perceptual experience, and we would not take it as such. We would take them as imaginings or hopings, or some other state. As conscious subjects we are aware that perceptual states present themselves to us as truth-connected – they purport to say how the world is. We also know that they are truth-connected since they can serve as evidence for perceptual belief – again, they have the status of evidence, in part, because perceptual belief can be justified on their basis.

Peacocke similarly notes in passing that perceptual experiences have specifically representational content, and not merely intentional content like hoping, wishing and visual imagining do, which in part explains how they can contribute to the positive epistemic status of beliefs. (Peacocke, 2004, p. 99) While all of these states have intentional content in that they are ‘about’ something, the difference between them and specifically representational states is this: “In being in a state with representational content the subject of the state is thereby under the impression that the world is a certain way.” (Ibid.) This is not so in the cases I mention of wishing and imagining. If, however, other states did represent their content assertively and we thought that their representing things a certain way did count towards their being true, they would be justificationaly relevant since they would be evidence.

In summary, perceptual experience is justificationaly relevant for perceptual belief in a way that other propositional states are not. Perceptual experience has a conceptual truth-connection in virtue of being evidence that these other states lack, and we are aware of this fact because of the unique way perceptual experience presents itself to the subject’s consciousness, and because we judge that beliefs formed on its basis are justified. The role of phenomenology here is to allow us to tell, from the first person

perspective, when we are dealing with truth-connected states (evidence). A truth-connected state, like perceptual experience, presents itself in a distinctive way to us in consciousness so that we have a way to distinguish it from non-truth connected states, which play no justificatory role.

We are now in a position to fully understand what the connections are between justified belief and truth. In the case of the justification of perceptual belief, justification supervenes on three things: belief, perceptual experience, and the relation between them.⁹⁵ A first aspect of justification's truth-connection is that it supervenes on truth-connected states. Justification is truth-connected in that the only states relevant to whether or not it obtains are conceptually tied to truth in virtue of their being evidence for what is believed. We know that such states are evidence for what is believed since belief can be justified on their basis, and we also know that only truth-related considerations can justify belief. Evidence, therefore, has a conceptual tie with truth. Not only can the truth-connection be seen in the generation of justification, but also in its loss. One loses justification if a defeater enters in at the level of belief, experience, or the basing relation. Recall that a defeater is something that provides one with a reason for ceasing to hold a belief, either by rebutting one's belief by being inconsistent with one's beliefs, or by undermining or undercutting one's reasons for thinking it *true*. The concept of truth is essentially connected both to how one acquires justification and to how one loses it.

A second aspect of the connection between justification and truth has specifically to do with the function of justification. Recall that, given the fundamentally first-personal nature of such concepts, I have argued that epistemic justification is at least in

⁹⁵ The necessary relation here must be one of the belief being "based" on the experience. For a discussion of the basing relation, see 1.3.

part a matter of possessing reasons, reasons to which the subject can appeal in wondering if her beliefs are true. (See chapter 1 and 2) The function of justification is to act as an indicator to the subject of reasons to think that what they believe is true in the form of evidence. One of the reasons we care about having justified beliefs is that we care about having beliefs that we take to be supported by good reasons, which are good because they count in favour of truth. If justification did not serve this truth-related function, it is unclear why we would care about justification or having epistemically justified beliefs.

As I argued above, a well functioning epistemic system is one that includes the subject having justified beliefs. The standards that govern whether or not epistemic justification obtain require that the subject have well-evidenced belief that meets the awareness condition I defended in Chapter 2.⁹⁶

The reason that some norms, like evidential norms such as if one is a state in which it appears to one that p , then believe that p (in the absence of defeaters), rather than others, like if one has a perceptual experience as of p , then believe that q (where q is some arbitrary proposition), is that we take ourselves to have good reason to think that some, but not all, putative norms are truth-connected. The only ones that are truth-connected are the evidential ones, which must be so, given what evidence is. An experience that p is evidence for the belief that p because (in the absence of defeaters), such a belief can be justified on that basis. This is different from believing q (where q is some arbitrary proposition) on the basis of an experience p , since q cannot be justified on

⁹⁶ By ‘norms’ of justification here I just mean standards, the meeting of which constitute a belief’s being epistemically justified. It does not seem that, *in general*, norms can only be applied if one believes that conditions are fulfilled, on pain of infinite regress. Rather, I have been attempting to make explicit conditions, the meeting of which, are necessary for epistemic justification to obtain. As most competent users of the relevant concepts that characterize our epistemic practice do not have beliefs about them (e.g. beliefs about the nature of epistemic justification), we must therefore *implicitly* grasp these concepts and standards and conduct ourselves in accordance with them. We can of course aim to make these concepts and standards *explicit* through philosophical reflection.

that basis. p is not evidence for q : if it were, then q would be justified (in the absence of defeaters) if properly based on p .

A pragmatic defense can also be given to vindicate evidential norms such as if one is in a state in which it appears to one that p , then believe that p (in the absence of defeaters), rather than believe that q on that basis (where q is some arbitrary proposition). Suppose that we allow that perceptual experiences as of p justify belief. Which belief(s) does it justify? Surely the answer is either nothing at all, or p . We judge it absurd to suggest that the experience justifies the subject in believing q , for some arbitrary q . So the choice is between believing nothing and the content represented in experience. But if the subject is not justified in believing anything then scepticism is true; but commonsense assures us that such radical scepticism is false. So if the subject is to engage in theory construction of any kind and think anything at all he must believe that p if he experiences that p (at least provisionally). Here the subject has good practical reasons to do cognitively whatever is necessary to engage in cognitive development of any sort.

A related aspect of the connection between justification and truth is that there is a very close connection between the *property* of justification and the *process* of justification, the latter of which is manifestly truth-directed. Robert Audi calls this the Process-Property Integration Thesis. (Audi, 1993b, p. 305) Roughly, the idea is that justification is the property of a belief that, when citing its justifiers can, “at least in principle, both show that it is justified and (conceptually) *constitute* justifying it.” (Ibid.) The aim (or function) of the *process* of justification is to show or provide reasons to think that what is believed is true. This is done by appealing to evidence (which, as a conceptual matter, is that which justifies belief). As Audi rightly notes, that is why it is

improper to claim to be interested in justifying belief but not be concerned with whether or not the belief is true. (Ibid., p. 319) The property of justification is truth-connected, therefore, in virtue of the closely associated process of justifying belief, which has truth as its intrinsic end.

Noting the close conceptual connection between the property and the process of justification, however, is not to confuse the *property* of justification with the *activity or process* of justification. These are related, but they are distinct. I have been suggesting that the Awareness Requirement demands that being aware of one's grounds is *necessary* for having justification. That is not to say that *showing* one's justification is *constitutive* of having justification. One might be tempted to think so, because in many cases when one is aware of something, one is often able to prove or show it. When the property of justification obtains, since one is aware of it, one is often able to express what one is aware of in conversation and form beliefs about it. Those utterances and beliefs are not what is doing the justifying, however; they are mere expressions of it.

When one has justification, it is necessary that one is aware of the belief's grounds, which is minimally to appreciate the existence, relevance and adequacy of one's grounds (see chapter 2). Again, it is not the process of justification that *makes* one justified. If, however, one is unable to engage in the activity of justifying, it may be relevant as to *why* one cannot. If one cannot show or prove one's justification, even to oneself, because one is *not aware* of one's grounds, then one is not justified.

A final issue needs to be addressed. While not suffering from blatant premise circularity, one might object that my account of epistemic justification's truth-connection in terms of evidence with its conceptual tie to truth, which I claim must be so, given the

nature of epistemic justification, is still in some sense objectionably circular. If by ‘circularity’ in this sense one means that one cannot understand what evidence is without understanding what epistemic justification is, and vice versa, then this may be true. However, I do not think that *this* form of circularity is necessarily objectionable. It might be a problem if I was intending a *reductive analysis* of the truth connection (or any concept whatsoever) in terms of necessary and sufficient conditions, since the point of a conceptual *reduction* is to explain the target concept in terms of others that make no reference to that target. But that has not been my aim.

My aim here, and throughout this thesis, has been to provide an *elucidation* of the concept of epistemic justification. Centrally, I have tried to cast light on a central feature of justification, namely a necessary condition of its obtaining, that I captured in the Awareness Requirement (see Chapter 2). In order to achieve a reflective understanding of the concept of epistemic justification, I have been employing a Strawsonian model of conceptual ‘analysis’ that aims for *connection* over *reduction*. Strawson understands that contrast as follows:

Let us imagine, instead, the model of an elaborate network, a system, of connected items, concepts, such that the function of each item, each concept, could, from the philosophical point of view, be properly understood only by grasping its connections with the others, its place in the system – perhaps better still, the picture of a set of interlocking systems of such a kind. If this becomes our model, then there will be no reason to be worried if, in the process of tracing connections from one point to another of the network, we find ourselves returning to, or passing through, our starting-point. We might find, for example, that we could not fully elucidate the concept of knowledge without reference to the concept of sense perception; and that we could not explain all the features of the concept of sense perception without reference to the concept of knowledge. But this might be an unworrying and unsurprising fact. (Strawson, 1992, p. 19)

So applying this to the case at hand, given how fundamental these concepts are, it need not be worrying and surprising that the concept of epistemic justification cannot be elucidated without reference to evidence; and that we cannot explain all the features of the concept of evidence, i.e. its conceptual connection to truth, without reference to the concept of epistemic justification.

As Strawson notes, that is not to say that certain kinds of circularity are *never* objectionable. Some circles are too small to establish a revealing connection, and we judge them as such. (Ibid., p. 20) What I hope to have done here, however, is drawn the connections between epistemic justification, evidence, and truth in a way that illuminates what makes epistemic justification epistemic.

In summary, I began this chapter by motivating and defending the idea that being justified is a kind of epistemic success or accomplishment distinct from knowledge. I then turned to the conceptual connections between justification and truth that individuate this kind of justification for other sorts.

By a conceptual connection I mean that, in virtue of the concepts involved, there are entailments between the concepts, and that grasping a concept is sufficient to grasp concepts that are conceptually connected to it. I argued that in virtue of the concepts themselves, epistemic justification and truth are conceptually connected. That is, certain entailments hold between the concepts of epistemic justification and truth, and that grasping the concept of epistemic justification is sufficient to appreciate these entailments *a priori*.

Specifically, I argued that epistemic justification is epistemic because it turns on evidence; evidence in turn is epistemic because it is conceptually linked with truth. That

is, it is not possible to competently employ the concept of evidence without appreciating that it counts in favour of the truth of that which it is evidence for, just as one cannot, for example, competently employ the concept of bachelorhood without appreciating that it applies to all and only unmarried men. Epistemic justification, therefore, is conceptually linked with the truth (via evidence), which is what makes it distinctively epistemic. This account of the truth-connection is substantial enough to ensure that the kind of justification we are concerned with is genuinely epistemic, while still being consistent with the kind of epistemic internalism I am advocating.

CONCLUSION:

In Chapter 1 I addressed some preliminary issues and defended some fundamental distinctions concerning the nature of epistemic justification. For instance, I outlined some of the basic features of justification, distinguishing doxastic from propositional justification, and considered how best to understand the epistemic basing relation. Specifically, I considered what role (if any) causation plays in properly basing a belief on a reason. I also considered the structure of justification and argued for the existence of immediate justification, that is, the justification that does not come from one's justification to believe other propositions. In effect I defended Foundationalism against its rivals (e.g., Coherentism). In addition I also explored ways of construing the internalism/externalism distinction, and considered to what kinds of epistemic evaluation the distinction applies. I concluded the chapter by outlining the egocentric or first person perspective and noting the kinds of constraints that such a perspective puts on an adequate theory of epistemic justification.

In Chapter 2 I motivated the need for *awareness* as a necessary condition of epistemic justification obtaining. I captured the kind of awareness justification requires with the so-called Awareness Requirement, which is the hallmark of the kind of internalism I am defending. The condition can be stated as follows:

Awareness Requirement:

S's is justified in believing that *p* only if

- ii) There is something, *X*, that contributes to the justification of S's believing that *p*;
and
- ii) For all *X* that contributes, S is aware (or potentially aware) of *X*.

I argued for this requirement by drawing attention to our considered judgements about cases. I considered two kinds of cases: first, I clarified and drew out the implications of cases of clairvoyance and other cases of unusual but reliable cognitive faculties. Secondly, I showed that even many of those who claim to reject such awareness requirements implicitly appeal to them to motivate their own accounts of justification. By considering the arguments for “Mentalism”, I showed that, unless an awareness requirement is presupposed, the cases adduced by Mentalists are of no intuitive force. In turn, by considering just these cases I tried to establish that the awareness requirement is necessary for justification.

In the latter half of Chapter 2 I provided an account of the nature of this awareness. In doing so I examined and responded to an argument against the awareness requirement that I called Bergmann’s Dilemma. My solution invoked what I call ‘non-doxastic strong awareness’ in a way that avoids the regresses that Bergmann advances. The overall aim of Chapter 2 was to establish that factors external to the subject’s consciousness (such as the reliability of the process that gives rise to the belief in question) are not *sufficient* for justification. Awareness is required.

In Chapter 3 I considered whether external factors are nevertheless *necessary* for justification to obtain. I argued that they are not. This was done by considering the case of the New Evil Demon. After having explained this case, I showed how it supports Fallibilism and the Awareness Requirement, as well as the claim that the *actual* reliability of the belief-forming mechanism used in a world is not necessary for epistemic justification to obtain. I then defended this position against recent influential arguments advanced by Timothy Williamson and John McDowell.

Some connection to truth is what makes a given thing epistemic. In Chapter 4 I investigated proposals that conceive of the relation between justified belief and truth as one of truth-conduciveness. I argued that, given internalism about justification, either in the sense I defend in Chapter 2 involving the Awareness Requirement, or in a weaker sense that Christopher Peacocke endorses, the truth-relation cannot be characterized in terms of truth-conduciveness or reliability. This is not just to reiterate a conclusion of Chapter 3 that justification need not be *actually* truth-conducive. I argued further that the truth-connection cannot even be one of truth-conduciveness in a *conditional* mode. Having dispatched truth-conducive accounts of the truth-connection, I developed a positive account of the truth-connection in Chapter 5.

Epistemic justification is different from other kinds of justification. Moral, aesthetic and pragmatic justification (assuming there are such things), seem to relate essentially to goodness, beauty and usefulness, respectively. What individuates *epistemic* justification from the other types is *some* connection to truth. Thus far we had seen what this connection is not: in Chapter 3 I argued that justification need not be *actually* truth-conducive. In Chapter 4 I argued further that the truth-connection cannot even be one of truth-conduciveness in a *conditional* mode.

In Chapter 5 I argued against recent work that assimilates justification to knowledge, as well as for a positive account of the truth connection. As to the former issue, I defended the orthodoxy that they are distinct epistemic statuses; as to the latter, I argued that the connection between justification and truth is *conceptual*. That is, epistemic justification is epistemic because it turns on evidence; evidence is epistemic because it is conceptually linked with truth. Epistemic justification, therefore, is

conceptually linked with the truth (via evidence), which is what makes it distinctively epistemic. This account of the truth-connection is substantial enough to ensure that the kind of justification we are concerned with is genuinely epistemic, while still being consistent with the kind of epistemic Internalism I am advocating.

Throughout this thesis I presupposed Particularism as a methodological constraint on approaching the theory of epistemic justification. That is, given that I share an intuitive idea of what we are justified in believing, I sought to give an account of what all the intuitive particular cases of justification have in common that make them instances of *justification*, rather than something else. To that end I argued for the indispensability of conscious awareness to epistemic justification.

With a clearer answer to the meta-epistemological question of what is fundamental about epistemic justification, I am now in a position to do some *applied* epistemology, viz., to consider what, if anything, we are justified in believing, and to ask how such justification is possible. Now that I have a better idea of what we are looking for, I can start to search for it, and explain its possibility. In the future I intend to extend the account developed here of the justification of perceptual belief to other potentially problematic areas, such as memory, testimony, and the *a priori*. In each case I intend to offer a response that explains how a type of justification is possible that is neither externalist, nor concessive to the sceptic, in a way that is consistent with epistemic internalism and the awareness condition.

In short, this thesis is an explanation and defense of an internalist theory of epistemic justification. The central claim of the thesis is that something is “internal” in this sense only if it is, or can easily be, the object of the agent’s conscious awareness. I

have argued throughout that conscious awareness is a necessary condition of epistemic justification obtaining, and that factors external to consciousness play no justificatory role.

BIBLIOGRAPHY:

- Alston, William. (1989a) 'Concepts of Epistemic Justification' reprinted in *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. (1989b) 'The Deontological Conception of Justification' reprinted in *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. (1989c) 'An Internalist Externalism' reprinted in *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. (1989d) 'Level Confusions in Epistemology' reprinted in *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. (1989e) 'Internalism and Externalism in Epistemology' reprinted in *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. (1993) *The Reliability of Sense Perception*. Ithaca: Cornell University Press.
- Armstrong, David. (1973) *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Austin, J.L. (1961) 'Other Minds' reprinted in *Philosophical Papers*. Oxford: Clarendon Press.
- Audi, Robert. (1982) 'Believing and Affirming' in *Mind* 91, pp.115-120.
- Audi, Robert. (1993a) 'Belief, Reason and Inference' in *The Structure of Justification*. Cambridge: Cambridge University Press.
- Audi, Robert. (1993b) 'Justification, Truth, and Reliability' in *The Structure of Justification*. Cambridge: Cambridge University Press.
- Audi, Robert. (2001) 'Doxastic Voluntarism and the Ethics of Belief' in *Knowledge, Truth, and Duty*. ed. M Steup. Oxford: Oxford University Press.
- Beilby, James, ed. (2002) *Naturalism Defeated?* Ithaca: Cornell University Press.
- Bergmann, Michael. (2006) *Justification Without Awareness*. Oxford: Clarendon Press.
- Bird, Alexander. (2007) 'Justified Judging' in *Philosophy and Phenomenological Research* 74, pp. 84-110.

- Bonjour, Laurence. (1980) 'Externalist Theories of Empirical Knowledge' in *Midwest Studies in Philosophy* 5, pp. 53-73.
- Bonjour, Laurence. (1985) *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.
- Bonjour, Laurence. (1998) *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Bonjour, Laurence. (2001) 'The Indispensability of Internalism' in *Philosophical Topics* 29, pp. 47-66.
- Burge, Tyler. (1979) 'Individualism and the Mental' reprinted in *Externalism and Self-Knowledge*, ed. Ludlow and Martin, Stanford: CSLI Publications, 1998.
- Burge, Tyler. (1982) 'Other Bodies' in *Thought and Object*, ed. A Woodfield. Oxford: Oxford University Press.
- Burge, Tyler. (2003) 'Perceptual Entitlement' in *Philosophy and Phenomenological Research* 67, pp. 503-548.
- Chisholm, Roderick. (1966) *Theory of Knowledge*. 1st ed. Englewood, N.J.: Prentice-Hall.
- Chisholm, Roderick. (1989) *Theory of Knowledge*. 3rd ed. Englewood, N.J.: Prentice-Hall.
- Coady, C.A.J. (1992) *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- Cohen, Stewart. (1984) 'Justification and Truth' in *Philosophical Studies* 46, pp. 279-295.
- Conee, Earl and Feldman, Richard. (1998) 'The Generality Problem for Reliabilism' in *Philosophical Studies* 89, pp. 1-29.
- Crane, Tim and Mellor, D.H. (1990) 'There is No Question of Physicalism' in *Mind* 99, pp. 185-206.
- Davidson, Donald. (1963) 'Actions, Reasons, and Causes' in *The Journal of Philosophy* 60, pp. 685-700.
- Davidson, Donald. (2000) 'A Coherence Theory of Truth and Knowledge' in *Epistemology: An Anthology*. ed. Sosa / Kim. Oxford: Blackwell Publishing.

- Feldman, Richard. (2001) 'Voluntary Belief and Epistemic Evaluation' in *Knowledge, Truth, and Duty*. ed. M Steup. Oxford: Oxford University Press.
- Feldman, Richard and Conee, Earl. (2001) 'Internalism Defended' in *Epistemology: Internalism and Externalism*. ed. Kornblith. Oxford: Blackwell Publishing.
- Feldman, Richard. (2002) 'Epistemological Duties' in *The Oxford Handbook of Epistemology*. ed. P Moser. Oxford: Oxford University Press.
- Feldman, Richard, Conee, Earl. (2004) 'Evidentialism' reprinted in *Evidentialism*. Oxford: Clarendon Press.
- Feldman, Richard. (2005) 'Justification Is Internal' in *Contemporary Debates in Epistemology*. ed. Steup and Sosa. Oxford: Blackwell Publishing.
- Foley, Richard. (1987) *The Theory of Epistemic Rationality*. Cambridge, Mass.: Harvard University Press.
- Foley, Richard. (1993) *Working Without a Net*. Oxford: Oxford University Press.
- Foley, Richard. (2005) 'Justified Belief as Responsible Belief' in *Contemporary Debates in Epistemology*. ed. Steup and Sosa. Oxford: Blackwell Publishing.
- Fumerton, Richard. (1995) *Metaepistemology and Skepticism*. Boston: Rowman and Littlefield.
- Gettier, Edmund. (1963) 'Is Justified True Belief Knowledge?' in *Analysis* 23, pp. 121-3.
- Ginet, Carl. (2001) 'Deciding to Believe' in *Knowledge, Truth, and Duty*. ed. M Steup. Oxford: Oxford University Press.
- Goldman, Alvin. (1979) 'What Is Justified Belief' in Pappas, G. ed. *Justification and Knowledge*. Dordrecht, Holland: D. Reidel.
- Goldman, Alvin. (1980) 'The Internalist Conception of Justification' in *Midwest Studies in Philosophy* 5, pp. 27-51.
- Goldman, Alvin. (1986) *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.
- Goldman, Alvin. (1992a) 'Discrimination and Perceptual Knowledge' reprinted in *Liaisons*. Cambridge, Mass: The MIT Press.

- Goldman, Alvin. (1992b) 'Strong and Weak Justification' reprinted in *Liaisons: Philosophy Meets The Cognitive and Social Sciences*. Cambridge, Mass.: The MIT Press.
- Goldman, Alvin. (1992c) 'A Causal Theory of Knowing' reprinted in *Liaisons*. Cambridge, Mass.: The MIT Press.
- Goldman, Alvin. (2001) 'Internalism Exposed' reprinted in *Epistemology: Internalism and Externalism*. ed. Hilary Kornblith. Oxford: Blackwell Publishing.
- Harman, Gilbert. (1986) *Change in View: Principles of Reasoning*. Cambridge, Mass.: MIT Press.
- Hetherington, Stephen. (2001) *Good Knowledge, Bad Knowledge*. Oxford: Clarendon Press.
- Jackson, Frank. (1998) *From Metaphysics To Ethics*. Oxford: Clarendon Press.
- Kaplan, Mark. (1996) *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Klein, Peter. (2000) 'A Definition Of Propositional Knowledge' reprinted in *Epistemology: An Anthology*. ed. Sosa / Kim. Oxford: Blackwell Publishing.
- Klein, Peter. (2005) 'Infinitism is the Solution to the Regress Problem' in *Contemporary Debates in Epistemology*. ed. Steup and Sosa. Oxford: Blackwell Publishing.
- Korcz, Keith Allen. (1997) 'Recent Work on the Basing Relation' in *The American Philosophical Quarterly* 34, pp. 171-191.
- Korcz, Keith Allen. (2000) 'The Causal-Doxastic Theory of the Basing Relation' in *Canadian Journal of Philosophy* 30, pp. 525-550.
- Lackey, Jennifer. (1999) 'Testimonial Knowledge and Transmission' in *The Philosophical Quarterly* 49, pp. 471-490.
- Lehrer, Keith and Paxon, Jr., Thomas. (1969) 'Knowledge: Undefeated Justified True Belief' in *The Journal Of Philosophy* 67, pp. 225-37.
- Lehrer, Keith. (1971) 'How Reasons Give Us Knowledge, or The Case of The Gypsy Lawyer' in *The Journal of Philosophy* 68, pp. 311-313.
- Lehrer, Keith. (1974) *Knowledge*. Oxford: Oxford University Press.

- Lehrer, Keith. (1990) *Theory Of Knowledge*. Boulder, CO: Westview.
- Lehrer, Keith. (2000) *Theory of Knowledge*, 2nd ed. Boulder, CO: Westview.
- Lenz, John. (1974) 'Problems for the Practicalist's Justification of Induction' reprinted in *The Justification of Induction*. ed. Richard Swinburne. Oxford: Oxford University Press.
- McDowell, John. (1996) *Mind and World*. Cambridge, Mass: Harvard University Press.
- McDowell, John. (1998a) 'Criteria, Defeasibility, and Knowledge' in *Meaning, Knowledge, and Reality*. Cambridge, Mass: Harvard University Press.
- McDowell, John. (1998b) 'Knowledge by Hearsay' in *Meaning, Knowledge, and Reality*. Cambridge, Mass: Harvard University Press.
- McDowell, John. (1998c) 'Knowledge and the Internal' in *Meaning, Knowledge and Reality*. Cambridge, Mass: Harvard University Press.
- McKim, Robert. (2001) *Religious Ambiguity and Religious Diversity*. Oxford: Oxford University Press.
- Moser, Paul. (1989) *Knowledge and Evidence*. Cambridge: Cambridge University Press.
- Neta, Ram and Pritchard, Duncan. (2007) 'McDowell and the New Evil Genius' in *Philosophy and Phenomenological Research* 74, pp. 381-396.
- Nozick, Robert. (1981) *Philosophical Explanations*. Oxford: Clarendon Press.
- Owens, David. (2003) 'Does Belief Have an Aim?' in *Philosophical Studies* 115, pp. 283-305.
- Peacocke, Christopher. (2004) *The Realm of Reason*. Oxford: Oxford University Press.
- Plantinga, Alvin. (1993) *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Plantinga, Alvin. (2000) *Warranted Christian Belief*. Oxford: Oxford University Press.
- Pritchard, Duncan. (2005) *Epistemic Luck*. Oxford: Clarendon Press.
- Pritchard, Duncan. (2006) *What Is This Thing Called Knowledge?* Abingdon: Routledge.

- Putnam, Hilary. (1973) 'Meaning and Reference' in *Journal of Philosophy* 70, pp. 699-711
- Pryor, James. (2001) 'Highlights of Recent Epistemology' in *British Journal for the Philosophy of Science* 52, pp. 95-124.
- Pryor, James. (2005) 'There is Immediate Justification' in *Contemporary Debates in Epistemology*. ed. Steup and Sosa. Oxford: Blackwell Publishing.
- Reichenbach, Hans. (1938) *Experience and Prediction*. Chicago: University of Chicago Press.
- Reichenbach, Hans. (1949) *The Theory of Probability*. Berkeley & Los Angeles: University of California Press.
- Salmon, Wesley. (1974) 'The Pragmatic Justification of Induction' reprinted in *The Justification of Induction*. ed. Richard Swinburne. Oxford: Oxford University Press.
- Scanlon, Thomas. (2000) *What We Owe to Each Other*. Cambridge, Mass: Harvard University Press.
- Schaffer, Jonathan. (manuscript) 'The Debasing Demon'.
- Shah, Nishi. (2003) 'How Truth Governs Belief' in *The Philosophical Review* 112, pp. 447-482.
- Shope, Robert K., ed. (1983) *The Analysis of Knowledge: A Decade of Research*. Princeton: Princeton University Press.
- Silins, Nicholas. (2005) 'Deception and Evidence' in *Philosophical Perspectives* 19, pp. 375-404.
- Smith, Nicholas. ed. (2002) *Reading McDowell: On Mind and World*. London: Routledge.
- Smithies, Declan. (manuscript) 'The Normative Role of Knowledge'.
- Snowdon, Paul. (2005) 'The Formulation of Disjunctivism: A Response to Fish' in *The Proceedings of the Aristotelian Society* 105, pp. 129-141.
- Sosa, Ernest. (2003) 'Beyond Internal Foundations to External Virtues' in *Epistemic Justification*. Oxford: Blackwell Publishing.
- Strawson, P. F. (1992) *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.

- Stroud, Barry. (2000) 'Transcendental Arguments' reprinted in *Understanding Human Knowledge*. Oxford: Oxford University Press.
- Sturgeon, Scott. (2000) *Matters of Mind*. London: Routledge.
- Sutton, Jonathan. (2005) 'Stick To What You Know' in *Nous* 39, pp. 359-396.
- Sutton, Jonathan. (2007) *Without Justification*. Cambridge, Mass.: The MIT Press.
- Swain, Marshall. (1979) 'Justification and the Basis of Belief' in Pappas, G. ed. *Justification and Knowledge*. Dordrecht, Holland: D. Reidel.
- Swain, Marshall. (1981) *Reasons and Knowledge*. Ithaca, NY: Cornell University Press.
- Swain, Marshall. (1985) 'Justification, Reasons and Reliability' in *Synthese* 64, pp. 69-92.
- Swinburne, Richard. (1995) 'Response to Warrant' in *Philosophy and Phenomenological Research* 55, pp. 415-419.
- Swinburne, Richard. (2001) *Epistemic Justification*. Oxford: Clarendon Press.
- Tolliver, Joseph. (1982) 'Basing Beliefs on Reasons' in *Grazer Philosophische Studien* 15, pp.149-161.
- Velleman, David. (2000) 'On the Aim of Belief' in *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- Wedgwood, Ralph. (2002) 'The Aim of Belief' in *Philosophical Perspectives* 16, pp. 267-297.
- Williams, Bernard. (1973) 'Deciding to Believe' in *Problems of the Self*. Cambridge: Cambridge University Press.
- Williamson, Timothy. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. (2007) 'On Being Justified in One's Head' in *Rationality and the Good*. ed. Timmons, Greco, Mele. Oxford: Oxford University Press.
- Zalabardo, Jose. (2006) 'Bonjour, Externalism and The Regress Problem' in *Synthese* 148, pp. 135-69.