
Theses and Dissertations

Summer 2012

Achieving epistemic descent

Brett Andrew Coppenger
University of Iowa

Copyright 2012 Brett Andrew Coppenger

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/3279>

Recommended Citation

Coppenger, Brett Andrew. "Achieving epistemic descent." PhD (Doctor of Philosophy) thesis, University of Iowa, 2012.
<https://ir.uiowa.edu/etd/3279>.

Follow this and additional works at: <https://ir.uiowa.edu/etd>

 Part of the [Philosophy Commons](#)

ACHIEVING EPISTEMIC DESCENT

by

Brett Andrew Coppenger

An Abstract

Of a thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Philosophy in the Graduate College of The University of Iowa

July 2012

Thesis Supervisor: Professor Richard Fumerton

ABSTRACT

Traditional accounts of justification can be characterized as trying to analyze justification in such a way that having a justified belief brings with it assurance of truth. The *internalist* offers a demanding requirement on justification: one's having a justified belief requires that one *see* what the belief has going for it. *Externalists* worry that the internalist's narrow conception of justification will lead to unacceptably radical and implausible skepticism. According to the externalist, one need *not* know what a belief has going for it in order for that belief to be justified. Externalism, though, comes with its own problems.

Ernest Sosa has attempted to bridge the divide between internalism and externalism by pairing the strengths of internalism (assurance) with the strengths of externalism (an answer to skepticism). Sosa distinguishes two kinds of knowledge: *animal knowledge* that is essentially externalist in character and *reflective knowledge* that is intended to capture our best intellectual procedure in regards to knowledge. On Sosa's view, one gains reflective knowledge by building upon (by adding further epistemic components to) animal knowledge. As a result, Sosa's view seems to illustrate a bottom-up approach to the analysis of knowledge (or justification): reflective knowledge is the result of animal knowledge and some other epistemic factor.

My project, in contrast to Sosa's, is to argue that one should start with an account of ideal justification (justification that is paradigmatically internalist) and then proceed by loosening the standards on ideal justification in an effort to develop the possibility of non-ideal kinds of justification. The view that I will develop will adopt Sosa's strategy of distinguishing kinds of knowledge (or justification), but will result in a top-down approach to

the analysis of justification. Instead of starting with an undemanding standard and layer levels on top, I will start with an ideal standard and strip layers away.

I will also argue that my view has some important advantages over Sosa's. Not only does Sosa's view seem to run into many of the problems that threaten externalism, but his view is incapable of offering the kind of assurance that the internalist is after. The view I develop will maintain the internalist's interest in assurance while also providing a response to some of the skeptical problems that have plagued internalists.

If my project is successful, then, even if the justification that results in many of the cases I will be exploring is (admittedly) not ideal, we can use these conceptions of justification to help explicate how one might have justified beliefs about a great number of things. The essentially internalist account that I have offered will not only illustrate a serious approach to dealing with skepticism, but it will also capture how many of our commonsensically justified beliefs are in fact justified (albeit in a less than ideal sense).

Abstract Approved:

Thesis Supervisor

Title and Department

Date

ACHIEVING EPISTEMIC DESCENT

by

Brett Andrew Coppenger

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy
degree in Philosophy in the Graduate College of The University of Iowa

July 2012

Thesis Supervisor: Professor Richard Fumerton

Copyright by

BRETT ANDREW COPPENGER

2012

All Rights Reserved

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Brett Andrew Coppenger

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Philosophy at the July 2012 graduation.

Thesis Committee:

Richard Fumerton, Thesis Supervisor

Evan Fales

Diane Jeske

David Cuning

Ali Hasan

For Megan

TABLE OF CONTENTS

CHAPTER

1.	MOTIVATING THE PROJECT	1
1.1	The Internalism/Externalism Controversies	1
1.1.1	<i>Internal-State</i> Internalism	1
1.1.2	<i>Awareness</i> Internalism	4
1.1.3	Internalism and Assurance of Truth	7
1.2	Goldman's Rejection of Internalism	9
1.2.1	Goldman's Presentation of Internalism	9
1.2.2	Problems of Unavailable Evidence	11
1.2.2.1	The Problem of Stored Beliefs	11
1.2.2.2	The Problem of Forgotten Evidence	12
1.2.3	Problems of Cognitive Limitations	14
1.2.3.1	The Problem of Concurrent Retrieval	14
1.2.3.2	The Problem of the Doxastic Interval	15
1.2.4	Goldman's Conclusion	15
1.3	The Challenge for Internalism	16
2.	PLACING THE PROJECT IN ITS CONTEXT	17
2.1	Sosa's Attempt at Achieving Epistemic Ascent	17
2.1.1	The Epistemically Undemanding Level	19
2.1.1.1	Results from the Undemanding Level	23
2.1.2	The Epistemically Demanding Level	24
2.1.2.1	The Problem of Easy Knowledge	27
2.1.2.2	Results from the Demanding Level	30
2.2	Sosa's Theory Applied	30
2.2.1	The New Evil-Demon Problem	30
2.2.1.1	Denying the Underlying Assumption	32
2.2.1.2	Strong and Weak Justification	32
2.2.1.3	Normal-World Reliabilism	33
2.2.1.4	Transglobal Reliabilism	34
2.2.1.5	Actual-World Reliabilism	36
2.2.2	Knowledge of the External World	37
3.	PRESERVING INTERNALISM	40
3.1	Motivating Internalism	40
3.1.1	The Normative Conception of Justified Belief	40
3.1.1.1	Rejecting the Normative Conception	42
3.1.2	Mentalism	45
3.1.2.1	Rejecting Mentalism	47
3.1.3	The Subject's Perspective Objection	49

3.2	Assurance Internalism	52
3.2.1	Motivating Assurance Internalism	53
3.2.2	Clarifications	54
3.2.3	The Case for Assurance Internalism	58
3.3	Acquaintance and Noninferential Justification	60
3.4	Acquaintance, Noninferential Justification, and Assurance	62
3.4.1	Objections and Replies	63
4.	ACHIEVING EPISTEMIC DESCENT	67
4.1	Rejecting Sosa's Account	67
4.1.1	The New Evil Demon Problem <i>Revisited</i>	67
4.1.2	The Problem of Easy Knowledge <i>Revisited</i>	71
4.2	The Apparatus of Epistemic Descent	74
4.2.1	Epistemic Descent and Noninferential Justification	76
4.2.2	Fumerton's Principle of Inferential Justification	82
4.2.3	Epistemic Descent and Inferential Justification	87
4.2.4	Clarifications	93
5.	THE DIVIDENDS OF EPISTEMIC DESCENT	97
5.1	Results from the Apparatus of Epistemic Descent	97
5.1.1	Goldman's "Internalism Exposed"	97
5.1.1.1	Problems of Unavailable Evidence <i>Revisited</i>	98
5.1.1.2	Problems of Cognitive Limitations <i>Revisited</i>	100
5.1.2	The Gettier Problem	101
5.1.3	Epistemic Descent vs. Epistemic Ascent	106
5.2	Is this a Satisfying Response to Skepticism?	111
	REFERENCES	121

CHAPTER 1

MOTIVATING THE PROJECT

1.1 The Internalism/Externalism Controversies

The internalism/externalism controversies about how to understand justification have taken center stage in contemporary epistemology. The competing conceptions of internalism (and externalism) employed by different philosophers (or by the same philosophers at different times) result in several quite different internalist/externalist controversies. In order to properly characterize these controversies a preliminary discussion of the relevant technical terminology is necessary.

1.1.1 *Internal-State Internalism*

One prominent characterization of internalism focuses on the internal states of the subject.¹ On this view the justification a person has for a belief at a certain time is constituted by the internal state of the believer at that time. “The internalist, on this view, maintains that ... having a justified belief that P, consists in S’s being in some internal state.”² If internal-state internalism is committed to understanding justification in terms of the internal properties of the subject, then the contrasting view, external-state externalism could be understood as the view that justification has something to do with external states.

¹ My presentation of internal-state internalism will closely follow that of Richard Fumerton. See his 1995a: 60–2. For a similar discussion see also Fumerton 2001 and Fumerton 2006b.

² Fumerton 1995a.

Internal-state internalism, in addition to making good sense of the locution of ‘internalism,’ also seems especially well-suited to capture the motivation behind a now-famous thought experiment. Consider the New Evil-Demon problem:³

Imagine two different cases; in the first case “unknown to us, our cognitive processes, those involved in perception, memory and inference, are rendered unreliable by the actions of a powerful demon or malevolent scientist.”⁴ In the second case the very same cognitive processes are not affected by a powerful demon; instead our processes work just as they should, reliably, in our environment. By hypothesis, in the two cases the background beliefs are the same, and the sensory experiences are phenomenally indistinguishable from one another.

According to the internal-state internalist the justificatory status of the beliefs of the subjects in the two cases should be the same as well. If the two subjects are in identical internal states, and one agent is justified, then by definition the other agent is justified.

However, according to external-state externalism, it is possible for two agents to be in the same internal state while the differing external states result in one of the agents to be justified and the other to be unjustified.

Paradigm presentations of this kind of view can be seen in Conee and Feldman’s articulation of mentalism:⁵ “If any two individuals are exactly alike mentally, then they are alike justificationaly, e.g., the same beliefs are justified for them to the same extent.” Because this way of carving up the internalism/externalism controversy turns on what counts as an internal state, any developed account of internal-state internalism would require an analysis of just what an internal (as opposed to external) state is supposed to

³ Reconstructed from Lehrer and Cohen 1983.

⁴ Lehrer and Cohen 1983: 192–3.

⁵ Conee and Feldman 2001.

be. The different views that could end up counting as varieties of internalism on this view may seem surprising.

One might worry, for instance, about the kind of internalism that would result given certain views on perception. According to the direct realist, for instance, the objects of awareness in veridical perception are the actual objects in the world. As a result, would the internal-state internalist want to make the objects in the world part of one's internal states? Another potential problem would result if the advocate of internal-state internalism has an affinity for externalism with regard to mental content. Again, it would seem that the combination of the two views would result in a very strange understanding of what counts as internal.

Even more worrisome would be the implications of adopting a version of externalism about mental content and its implications for the type of internal-state internalism that would result. Such a view, while internalist by definition, surely seems at odds with the intended distinction being made.

These considerations seem to bring to the fore that the internal-state internalist needs to very carefully (and non-dogmatically) characterize internal states in a way that does not bring too much in. Of course, the problem is made even more difficult by the realization that not only must the internal-state internalist be careful about what she lets into her internalism, but she must not exclude too much either. It would seem problematic, for instance, if one were to limit internal-state internalism to non-relational properties of mind if the cost of such a distinction was a lack of justified beliefs about necessary truths.

Another requirement that the internal-state internalist must satisfy is to provide an analysis of which internal states result in justification. Even if it is granted that internal states are the only relevant feature in the evaluation of justification, work still must be done in articulating why certain internal states result in justification while others do not. For example, in a case of veridical perception, what is it about a subject's internal states that confer justification?

It is worth emphasizing that even if these two requirements for the internal-state internalist are met, one might wonder whether this way of carving up the difference between internalism and externalism gets to the heart of the issue.

1.1.2 Awareness Internalism⁶

Another prominent characterization of internalism focuses on notions of awareness.⁷ On this view the justification a subject has for a belief at a certain time is dependent on the subject's being aware of what a belief has going for it. S's belief B is justified only if (i) there is something, X, that contributes to the justification of B – e.g. evidence for B, a truth indicator for B, or the satisfaction of some necessary condition of B's justification – and (ii) S is aware (or potentially aware) of X.⁸ On this way of understanding justification the internalist is committed to a requirement of awareness of the justification contributor. The contrasting view, awareness externalism, could be

⁶ What I am calling awareness internalism is intended to be identical to what others have called *access* internalism. See BonJour 2002, Fumerton 1995a and 2006a, and Conee and Feldman 2004.

⁷ My presentation of awareness internalism will closely follow that of Michael Bergmann. See Bergmann 2006: 9–11.

⁸ Bergmann 2006: 9.

understood as the idea that beliefs can be justified even if the subject has no awareness of the justification contributor of her belief.

Just as internal-state internalism was, awareness internalism can be motivated by way of a thought experiment that seems to characterize an essential difference between internalism and externalism. Consider the case of Norman the clairvoyant:

Norman is a clairvoyant, and he has the ability to reliably form beliefs as a result of this power. Norman has no evidence of any kind for or against the possibility of this kind of power, nor does he have any reason to believe or not believe that he has such a power. One day, Norman comes to believe that the President is in New York City, though he has no evidence for or against this belief. Unbeknownst to Norman, the belief was a result of his reliable clairvoyant power, and the belief is in fact true.⁹

According to the awareness internalist the reason why Norman lacks justification when it comes to his belief about the President's location is because he has no awareness of something that would justify the belief. Even though the belief is true (and reliably produced), from Norman's perspective the belief is on par with a hunch or wishful thinking.

With cases like this in mind, BonJour goes on to articulate the difference between the awareness internalist and externalist:

The fundamental claim of internalism ... is that epistemological issues arise and must be dealt with from within the individual person's first-person cognitive perspective, appealing only to things that are accessible to that individual from that standpoint.¹⁰

An adequate defense of awareness internalism would offer an analysis of the particular kind of awareness that is required in order to have a justified belief. Varieties of awareness internalism differ according to whether they require that a subject actually

⁹ Reconstructed from BonJour 1980: 53–73.

¹⁰ BonJour 2010: 222.

have awareness of the justification contributor or, alternatively, if the awareness required only amounts to potential awareness to the justification contributor.

Thus, for example, Conee and Feldman develop the idea of something less than actual awareness when they argue that “some non-occurrent [mental] states that one is already in, such as non-occurrent memories of perceptual experiences, are stored evidence. Presently having this stored evidence justifies dispositionally some non-occurrent beliefs that one already has.”¹¹

The awareness internalist must also specify the nature of the (actual or potential) awareness that is required. One could construe the required awareness as being conceptual in nature – something that involves an act of judgment. Alternatively, one might contend that the required awareness need not involve an act of judgment. The former view can be construed as strong awareness, according to which a subject *conceives* of the justification contributor that is the object of awareness as being in some way relevant to the justification or truth of the belief. The latter view can be labeled weak awareness, according to which the required awareness does not involve conceiving of the justification contributor as relevant to the truth or justification of the belief.

Even if the awareness internalist clearly defines the kind of awareness that is required for justification (actual or potential, conceptual or non-conceptual), she must still clearly articulate the nature of the justification contributor that is required for justification. In other words, not only is awareness required for justification, but one must have awareness of a certain kind of thing in order to be justified. Again, the varieties of

¹¹ Conee and Feldman 2001: 8.

awareness internalism that result from the many different possible ways of developing the justification contributor make for very different accounts of internalism.

One could argue that justification requires awareness of each of the factors contributing to one's justification. On this view, again, it would seem like the internalist is committed to bringing too much into her internalism – all of the causal facts that led to the belief. Or one could ease the restrictions on justification and contend that justification requires awareness of something that entails (or, more weakly, makes probable) the proposition believed.

Perhaps then the internalist should require awareness of some X and awareness of X's entailing or making probable the proposition believed. Again, though, even if the requirements on a developed version of awareness internalism can be met, one might still wonder whether this way of carving up the internalism/externalism debate gets to the heart of the issue.

1.1.3 Internalism and Assurance of Truth

The characterization of internalism that I am most interested in, and the one that I will argue (in Chapter 3) best captures the heart of the internalist/externalist controversy relies on the idea that having a justified belief brings with it assurance of truth. The internalist aims at capturing this sense of assurance by offering a demanding account of justification where assurance of truth results in justified belief. When one has a justified belief, she will be in a position to *see* what the belief has going for it. On this conception of internalism, the fundamental feature of justification should be the assurance of truth that accompanies justified belief. The contrasting view to internalism that brings with it assurance of truth is a version of externalism according to which justified belief brings

with it no assurance of truth. On this way of carving up the debate between internalism and externalism, it should be obvious that according to the externalist a person can satisfy the required conditions on justification, and be justified in a belief, without having any idea that the belief is true. In essence, the internalist is committed to the idea that justification requires seeing from the first-person perspective what a belief has going for it, or, put another way, having some sense of assurance that a belief is true. The externalist rejects the idea that justified belief will result in assurance of truth.

It is my contention that this characterization of internalism is importantly different from both internal-state internalism and awareness internalism. First, in regard to internal-state internalism, by making explicit internalism's preoccupation with assurance of truth, the issue of immediate importance is skepticism. I will argue that it is the internalist's attempt to offer assurance of truth that uniquely distinguishes internalism from externalism. Second, in regard to awareness internalism, by putting a further condition of awareness on justified belief I believe that the awareness internalist is inadvertently committing herself to a vicious regress. On my view, the assurance of truth that I am interested in results from the analysis of justified belief, and as a result, the regress problem that threatens the awareness internalist does not threaten the kind of internalism that I will argue for. Because a fundamental feature of the kind of internalism I will defend is its preoccupation with assurance of truth, the now-common argument that internalism results in unpalatable skepticism must first be considered.

1.2 Goldman's Rejection of Internalism

In his article "Internalism Exposed,"¹² Alvin Goldman argues that internalism encounters unacceptable and implausible skepticism. Goldman aims to challenge the "general architecture of internalism, and the attempt to justify this architecture by appeal to a certain conception of what justification consists in."¹³ Goldman contends that the skeptical problems that threaten all forms of internalism motivate the rejection of internalism in favor of externalism.

1.2.1 Goldman's Presentation of Internalism

In an effort to capture the rationale behind many versions of internalism Goldman proposes the following argument:

- (1) The *guidance-deontological (GD) conception of justification* is posited.
- (2) A certain constraint on the determiners of justification is derived from the GD conception, that is, the constraint that all justification determiners must be *accessible to*, or *knowable by*, the epistemic agent.
- (3) The accessibility or knowability constraint is taken to imply that only internal conditions qualify as legitimate determiners of justification. So justification must be purely an internal affair.¹⁴

Goldman finds historical support for his presentation of internalism as being closely connected to a guidance-deontological conception of justification. On this view, justification is taken to be a normative concept. Goldman understands guidance-deontological conceptions of justification as being constituted by two closely related claims. The first claim, the deontological conception of justification, is that justified

¹² Goldman 1999: 271–93.

¹³ Goldman 1999: 272.

¹⁴ Goldman 1999: 272.

belief in a proposition amounts to a person's being permitted (or obligated) to believe that proposition, while unjustified belief in a proposition means that belief of the proposition is not permitted. The second component, the guidance component, is that a person should guide her belief-formation so as to satisfy her epistemic duty.

The adoption of the GD conception of justification seems to motivate constraints on determiners of justification. Since an adherent to the GD conception of justification will be interested in determining for which propositions belief is permitted, it is necessary to restrict possible determiners of justification to the features of belief that are accessible to that person. "If you cannot accurately ascertain your epistemic duty at a given time," Goldman asks, "how can you be expected to execute that duty, and how can you reasonably be held responsible for executing that duty?"¹⁵

Goldman demonstrates that if internalists are committed to the claim that determiners of justification need to be restricted to what is accessible, then they must hold that the only facts that are accessible are justifiers. As a result Goldman presents what he calls the *knowability constraint on justifiers* (KJ):

KJ: The only facts that qualify as justifiers of an agent's believing P at time *t* are facts that the agent can readily know, at *t*, to obtain or not to obtain.¹⁶

Goldman proceeds by arguing that because KJ is fundamental to internalism, and since the KJ is threatened by a host of skeptical problems, one should not hold out hope for internalism. I will present four different problems that Goldman argues cause trouble for internalism by way of causing trouble for different conceptions of the KJ. The problems

¹⁵ Goldman 1999: 274.

¹⁶ Goldman 1999: 274.

can be grouped into two different categories: first, those that deal with unavailable evidence, and second, those that deal with cognitive limitations.¹⁷

1.2.2 Problems of Unavailable Evidence

1.2.2.1 The Problem of Stored Beliefs

Goldman argues that the vast majority of the beliefs we take to be justified require justification contributors that are simply beyond what the internalist can offer. Goldman first recognizes that the vast majority of the beliefs we commonsensically consider to be justified are not occurrent. The majority of our beliefs, beliefs that most people take to be justified, are stored beliefs. Goldman goes on to argue that there is nothing in one's present conscious state that seems well-suited to serve as a justification contributor for the stored beliefs:

At any given time, the vast majority of one's beliefs are stored in memory rather than occurrent or active. ...Furthermore, for almost any of these beliefs, one's conscious state at the time includes nothing that justifies it. No perceptual experience, no conscious memory event, and no premises consciously entertained at the selected moment will be justificationaly sufficient for such a belief.¹⁸

As a result, it would seem like the internalist, in spite of the fact that we commonsensically think many of our stored beliefs are justified, must be committed to admitting that non-occurrent or non-active (stored) beliefs are not justified. However, such a conclusion seems to entail a rather radical version of skepticism: justified belief, according to internalism, is at least limited to beliefs that are occurrent or active (non-stored).

¹⁷ Goldman's presentation of the problems also includes related problems that I will not discuss due to considerations of brevity. However, I will present what I take to be the most fundamental and challenging problems to internalism.

¹⁸ Goldman 1999: 278

1.2.2.2 The Problem of Forgotten Evidence

Goldman proceeds by arguing that even if one only considers occurrent beliefs, the internalist encounters further skepticism-breeding consequences. It is easy to imagine plenty of cases of occurrent belief where the evidence, or justification contributor, that originally led to the formation of the belief is no longer available. I believe that Eli Whitney invented the cotton gin. I do not have any good idea when I learned this, and only a best guess as to where and how: during one of my elementary school years by either reading it in a textbook or hearing it from my teacher (though I do not remember who my teacher was). It seems like many of my beliefs, even the occurrent beliefs that I presently entertain, are beliefs that are epistemically similar to my belief about Eli Whitney.

Goldman argues that even if a non-occurrent belief *was* justified according to the internalist standards, it is no longer clear how the internalist can accommodate the claim that the belief *is still* justified even if the justification contributor is no longer available:

Many justified beliefs are ones for which an agent once had adequate evidence that she subsequently forgot. At the time of epistemic appraisal she no longer possesses adequate evidence that is retrievable from memory.¹⁹

To illustrate the problem, Goldman considers a case regarding the health benefits of eating broccoli:

Last year, Sally read a story about the health benefits of broccoli in the “Science” section of the *New York Times*. She then justifiably formed a belief in broccoli’s beneficial effects. She still retains this belief but no longer recalls her original evidential source.²⁰

¹⁹ Goldman 1999: 280.

²⁰ Goldman 1999: 280.

The internalist, Goldman argues, could still attempt to account for the justification of Sally's belief by appeal to background beliefs that Sally might have. Perhaps Sally believes that most of what she remembers was learned in an epistemically proper way.²¹ Given this background belief, shouldn't we conclude that Sally's belief is justified?

Goldman argues that this approach runs into trouble:

In a variant case, suppose that Sally still has the same background belief – namely, that most of what she remembers was learned in an epistemically proper manner – but she in fact acquired her broccoli belief from the *National Enquirer* rather than the *New York Times*.²²

Goldman argues that in the *National Enquirer* case Sally's belief is not justified. The problem with internalism, as he sees it, is that the internalist cannot account for the relevant difference in Sally's belief formation between the two different cases. On his view, it is the causal origin of the beliefs that are relevant to determining the beliefs' justificatory status. But, as Goldman points out, "All past events [like causal origin] are 'external' and therefore irrelevant according to internalism."²³

Goldman's case shows that the internalist, even when considering occurrent belief, seems ill-suited when it comes to trying to make sense of how many of the beliefs we commonly hold to be justified are in fact justified. Since Sally does not remember what justified her belief about broccoli, how could such a belief be justified? After all, she could be in exactly the same mental state, or have exactly the same evidence available, as she would in a case in which her belief is unjustified.

²¹ Goldman 1999: 280.

²² Goldman 1999: 280.

²³ Goldman 1999: 280.

1.2.3 Problems of Cognitive Limitations

1.2.3.1 The Problem of Concurrent Retrieval

In addition to the problems that threaten internalism that result from the alleged lack of evidence required to justify belief, Goldman points out problems that result from our limitations as cognitive agents. Even if it is granted that “only conscious and stored mental states are justifiers,” this does not imply that “all sets or conjunctions of such states qualify as justifiers.”²⁴ Goldman argues that our cognitive limitations restrict the available sets of conscious and stored mental states beyond what is useful for internalists.

According to coherentism, justification requires, at least, consistency amongst one’s beliefs. However, consistency is not enough – one must be aware that her beliefs are consistent. As a result, the coherentist would need to concurrently entertain all of her beliefs, “but such concurrent retrieval is psychologically impossible.”²⁵ Goldman also argues that foundationalist theories do not fare much better. “Internalist foundationalism might also require concurrent retrieval of more basic (or low-level) beliefs than it is psychologically feasible to retrieve.”²⁶

Again, if Goldman is right, the internalist seems to face severe skeptical consequences. Because of the limitations on what is concurrently cognitively accessible to a person, the internalist seems incapable of making sense of how most of our commonsensically justified beliefs are in fact justified.

²⁴ Goldman 1999: 281.

²⁵ Goldman 1999: 282.

²⁶ Goldman 1999: 282.

1.2.3.2 The Problem of the Doxastic Decision Interval

Goldman argues that further problems result given our cognitive limitations.

Goldman realizes that once the internalist places a knowability constraint on justification, she should worry about the time the required cognitive operation will take:

If justification is contingent on the agent's ability to know what justifiers obtain, the agent should not be permitted to believe a proposition p at t unless she can know *by* t whether the relevant justifiers obtain. Since it necessarily takes some time to compute logical or probabilistic relations the simultaneity model of justification needs to be revised so that an agent's mental states at t justify her in believing only p at $t + \epsilon$, for some suitable ϵ . The value of ϵ cannot be too large, of course, lest the agent's mental states change so as to affect the justificational status of p .²⁷

I take the thrust of Goldman's argument to be that the internalist will once again need to limit the class of justification contributors. If, for instance, it takes someone so long to work through a proof that by the time she reaches the end of the proof she has let the initial premises slip out of her mind, then it seems odd to count those initial premises as part of her evidence for the proof. As a result, the internalist seems committed to the idea that computations must be quick and easy enough that one can *see* the conclusion and its relation to the premises. However, it would seem that this restriction entails further skeptical consequences.

1.2.4 Goldman's Conclusion

Goldman takes himself to have shown that internalism, regardless of which variation is endorsed, "does not survive the glare of the spotlight."²⁸ By arguing that the fundamental feature of internalism is the knowability constraint and showing that holding it results in unpalatable skepticism, Goldman contends that internalism should be rejected

²⁷ Goldman 1999: 283–4.

²⁸ Goldman 1999: 293.

in favor of a version of externalism that does not suffer from the same skeptical consequences.

1.3 The Challenge for Internalism

To his credit, I think that Goldman is largely right about the skeptical consequences of internalism. Perhaps, though, the skeptical consequences of internalism should be stated more carefully. It is true, I think, that according to certain epistemically demanding versions internalism, many of our beliefs are unjustified. However, it is my contention, contra Goldman, that the skeptical consequences of such a version of internalism do not constitute a defeater for the view.

I will argue that internalism, when developed in the right way, has the resources to dilute to a great degree the skeptical consequences that seem to result from the view. In what follows I will argue for the claim that internalism offers the most philosophically interesting account of justification (Chapter 3). I will also argue that by starting with an internalist account of ideal justification, one can develop degenerate kinds of justification by relaxing the constraints on justification (Chapter 4). These degenerate kinds of justification will allow the internalist to answer Goldman's skeptical challenge (Chapter 5), and, in the end, show why internalism has not been "exposed."

CHAPTER 2

PLACING THE PROJECT IN ITS CONTEXT

2.1 Sosa's Attempt at Achieving Epistemic Ascent

In his bid to respond to different skeptical problems, Ernest Sosa has famously distinguished between kinds of knowledge. One particularly interesting result of Sosa's epistemology is the way his analysis of knowledge seems to bridge two common divides: first, the divide between internalists and externalists; second, between foundationalists and coherentists. On Sosa's view, one can start with an epistemologically undemanding concept of knowledge, and then, by adding further epistemic elements, work her way up to a more philosophically satisfying kind of knowledge. By "climbing the epistemic ladder" from a basic kind of knowledge to a more satisfying form of knowledge we achieve epistemic ascent.

Sosa construes his account of knowledge as a kind of virtue epistemology, according to which "knowledge is belief whose success is 'creditable' to the believer."²⁹ Sosa treats belief as a kind of performance – something that can be evaluated. Belief, when aimed at truth, can achieve different levels of success. Sosa's enduring illustration makes clear the different ways in which a belief can succeed:

The archer's shot is a good example. The shot aims to hit the target, and its success can be judged by whether it does so or not, by its accuracy. However accurate it may be, there is a further dimension of evaluation: namely, how skillful the shot is, how much skill it manifests, how adroit it is. A shot might hit the bull's-eye, however, and might even manifest great skill, while failing utterly, as a shot, on a further dimension. Consider a shot diverted by a great gust of wind initially so that it would miss the target altogether but for a second gust that puts it back on track to hit the bull's-eye. This shot is both accurate and adroit, yet it is not accurate because it is adroit, so as to manifest the archer's skill and competence. It

²⁹ Sosa 2011: 86.

thus fails on a third dimension of evaluation, besides those of accuracy and adroitness: it fails to be apt.³⁰

Sosa's preferred illustration is not unique. The same analysis of belief can be seen with a number of different examples. Imagine a basketball player's shot. The shot may be accurate. It might succeed by falling through the basket. Of course, not everyone's ability to shoot a basketball is equal. If Kobe Bryant and I both took a shot from the same place on the court, Kobe's shot would be better than mine. Kobe's shot is the result of a skill-set that I do not possess. His shot was adroit, while mine was not. We can further evaluate Kobe's shot by noticing that the shot's accuracy may be the result of the appropriate skill-set. If the shot was accurate because it was adroit, then the shot was apt.

The usefulness of these examples is that they show us belief can be successful at different levels. Belief is successful at one level if it is *accurate*, at a higher level if it is *adroit* (reliable or competent), and at a still higher level if it is *apt* (accurate because of its reliability or competence). It is by means of these different levels of success that Sosa develops his AAA (accurate, adroit, apt) conception of knowledge. Critical to a defense of the to the AAA conception is Sosa's distinction between *animal* knowledge and *reflective* knowledge.

Before proceeding, further consideration of Sosa's characterization of aptness, which, as we will see, plays a crucial role in his account of knowledge, seems important. There is certainly something intuitive about the kinds of examples that Sosa uses to illustrate aptness. Surely, it is argued, there is an important sense in which the skilled archer's shot is successful according to whether the action manifests the archer's skill. Yet, we should not quickly assume that this intuitive idea is capable of a precise

³⁰ Sosa 2011: 4.

explication. Consider the case in which the archer's shot fails to achieve aptness: a shot that hits the bull's-eye but does so as a result of being diverted by wind. According to Sosa, such a shot is not apt.

It is not clear to me why this shot fails to be apt. The arrow did succeed in hitting the target, and the arrow would not have done so if not for the skill of the archer, wind or no wind. Perhaps though, Sosa could appeal to the archer's luck in order to show why the shot does not count as apt belief. However, wouldn't Sosa also want to say that a successful hit by a batter in Major League Baseball exemplifies an instance of the manifestation of the batter's skill and competence? Yet, the likelihood of a hit for even the best batter in baseball is more improbable than it is probable.³¹ Put another way, the manifestation of skill often *does* involve a certain amount of luck. But if this is the case, then one might worry about the way that Sosa dismisses the aptness of the archer's shot just because of a fortuitous breeze.

2.1.1 The Epistemically Undemanding Level

On Sosa's view, knowledge is more than mere true belief. Beliefs may be accurate, but sometimes only because the believer was lucky. He argues that the first level of epistemic evaluation worthy of the label "knowledge" requires that a belief's accuracy be the result of a reliable process. He develops his epistemically undemanding notion of animal knowledge as requiring *apt* belief:

One has animal knowledge about one's environment, one's past, and one's own experience if one's judgments and beliefs about these are direct responses to their impact – e.g. through perception or memory – with little or no benefit of reflection or understanding.³²

³¹ For a discussion of this kind of case, see Greco 2010: 76–80.

³² Sosa 2001: 240.

[Animal knowledge requires] a true belief whose justification by its source in intellectual virtue is *prima facie* but not overridden.³³

... “Animal” knowledge [requires] apt belief *without* requiring *defensibly* apt belief, i.e., apt belief that the subject aptly believes to be apt, and whose aptness the subject can therefore defend against relevant skeptical doubts...³⁴

... Animal knowledge does not require that the knower have an epistemic perspective on his belief, a perspective from which he endorses the source of that belief, from which he can see that source as reliably truth conducive.³⁵

Animal competence does not require the believer to endorse the reliability of the competence; *nor* does it require the believer to endorse the appropriateness of the conditions for the exercise of the competence in forming that belief.³⁶

It is clear from these characterizations of animal knowledge that, at the undemanding level, Sosa is developing a version of reliabilism, and like other versions of reliabilism, it is an externalist conception of knowledge (or justification).³⁷ It would, however, be uncharitable to lump Sosa’s account in with other generic versions of reliabilism – the view is too sophisticated for such a rash treatment. Perhaps the best way

³³ Sosa 2001: 240–1.

³⁴ Sosa 2007: 24.

³⁵ Sosa 2009: 135.

³⁶ Sosa 2011: 149–50.

³⁷ Sosa seems to slip between thinking that justification is limited to the higher levels of knowledge (“Animal knowledge will generally be apt belief but rarely if ever justified” (2001: 290)) and allowing for the existence of “animal justification” as opposed to some kind of higher level justification (“At an unreflective level, epistemic justification can hence derive from the holding of a condition whose absence is no more subjectively distinguishable from its presence than is a realistic dream from waking life. Still, without reflective, non-arbitrary assurance that you satisfy that condition, you cannot know reflectively something you might still know at the animal level” (2007: 16)).

to understand the intricacies of Sosa's view is to contrast it with a generic version of process reliabilism.

According to process reliabilism, justification is essentially a matter of a belief's having the right kind of history. In a simplistic presentation, the process reliabilist's view can be articulated in the following way: a belief is justified iff the belief is the result of a reliable process.

While the process reliabilist restricts justified belief to belief that is the result of a reliable process, Sosa limits his version of *virtue reliabilism* to beliefs that result from the agent's competencies, where he understands a "competence" as being a certain set of abilities.³⁸ From the agent's perspective, one need not be aware that the belief is the result of a competence in order to have animal knowledge. Here is Sosa's explication:

A performance is apt if its success manifests a competence seated in the agent (in relevantly appropriate conditions). It does not matter how fragile was the competence, or its appropriate conditions, when the agent issued the performance.³⁹

We can see from Sosa's dismissal of a limiting condition of "fragility" that animal knowledge lacks a safety requirement.⁴⁰ Because animal knowledge only requires apt belief, one can have animal knowledge even if it is true that the belief could very easily have been wrong. To see this, consider the now familiar Fake-Barn County Case:

Suppose we are told that, unknown to Henry, the district he has just entered is full of papier-mache facsimiles of barns. These facsimiles look from the road exactly like barns, but are really just facades, without back walls or interiors, quite incapable of being used as barns. They are so

³⁸ Sosa 2011: 80.

³⁹ Sosa 2011: 7.

⁴⁰ For our present purposes we can define the kind of safety that Sosa rejects as the requirement that S would not believe that P without it being true that P.

cleverly constructed that travelers invariably mistake them for barns. Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the object on that site were a facsimile, Henry would mistake it for a barn.⁴¹

The subject in Fake-Barn County may have animal knowledge that the thing before him is a barn, even if as it turns out, he is very lucky to be picking out the only *real* barn in the county.⁴²

Sosa also helpfully illustrates his account of animal knowledge using the example of color perception. Imagine seeing that a surface appears red in normal light, where the lighting could very easily have been bad. According to Sosa, “so long as the light is good ... you can manifest your fine color eyesight in believing the surface to be red.”⁴³ He goes on to argue that this belief manifests the competence of the agent even if, unbeknownst to the agent, the surface could very easily have been white but made to appear red by a red light.

There is one more feature of Sosa’s virtue reliabilism that is worth pointing out here. Like other varieties of reliabilism, knowledge (or justification) can be construed as foundational in character. As Goldman famously pointed out, it can be helpful in articulating an account of justification to distinguish between a base clause and a recursive clause of justification.⁴⁴ The base clause on this recursive analysis of

⁴¹ Goldman 1976.

⁴² It is at least curious that Sosa allows for animal knowledge in the kind of Gettier case just discussed. Recall Sosa’s contention that an archer’s shot that hits the bull’s-eye as a result of a fortuitous gust of wind is not even apt. This curiosity will be discussed at more detail in Chapter 4.

⁴³ Sosa 2011: 25.

⁴⁴ Goldman 1979.

justification identifies how noninferentially justified beliefs are in fact justified. This base clause is then used to identify how inferentially justified beliefs might also be justified. Again following Goldman, one would be wise to point out that for inferential justification, the belief-dependent process at play should be understood as being a conditionally reliable process. If the beliefs that serve as inputs are unjustified, one would not expect justified beliefs to result – junk in, junk out.

One can imagine a belief that is the direct result of interaction with one's environment that, in accordance with Sosa's view, then rises to the level of animal knowledge in virtue of being aptly held. One can also imagine these beliefs serving as inputs for a competence that has as outputs other beliefs that also rise to the level of animal knowledge. Importantly, according to Sosa, neither of these competencies needs to be within the cognitive grasp of the agent in order to count as animal knowledge.

From the above considerations it is clear that Sosa's sophisticated version of reliabilism is able to cash out at least two important contemporary themes in epistemology: his account of animal knowledge is properly construed as being both externalist and foundational in character.

2.1.1.1 Results from the Undemanding Level

Sosa is able to use his account of animal knowledge to make sense of the externalist's intuition that internalists have over-intellectualized the conception of knowledge (and justification) in terms of what is required in order to know.⁴⁵ One

⁴⁵ Again, Sosa seems to vacillate between treating justification as parallel to his animal/reflective knowledge distinction and treating justification as something that only applies to reflective knowledge. It is at least likely that Sosa would be inclined to say that this first level concept is a kind of knowledge. After all, Sosa might say, when we talk

attractive feature of Sosa's view is its ability to make sense of many of the beliefs that we commonly take to be justified.

It is important to realize just how minimal a standard Sosa has in mind when he is developing his account of animal knowledge. Sosa has used the example of a supermarket door (or automatic door) to illustrate just how basic animal knowledge really is. The automatic door is programmed to open when it senses a person approaching. Thus, we can attribute animal knowledge to an automatic door when it reacts appropriately to its environment.

The epistemically undemanding animal knowledge offers Sosa the resources he needs to avoid the skeptical problems that seem to plague the internalist. Surely (it is argued) young children, or even the higher animals, can know things, yet they would fail to meet the requirements on justification set forth by internalists. Sosa need only point out that in many cases the higher animals, cognitively challenged humans, or the rest of us (when not thinking reflectively) can still *know* – we have (animal) knowledge as long as our belief is apt.

2.1.2 The Epistemically Demanding Level

Sosa argues that a satisfying account of knowledge must move beyond the unreflective level. Sosa distinguishes his inherently externalist conception of animal knowledge from a more philosophically satisfying internalist account of reflective knowledge. Sosa's conception of reflective knowledge can be understood as requiring the addition of further epistemic elements to those that were required for animal knowledge:

about dogs and fish we *naturally* talk about their having *knowledge*, but it at least sounds odd to talk about these animals as having *justified* belief.

One has reflective knowledge if one's judgment or belief manifests not only such direct response to the fact known but also understanding of its place in a wider whole that includes one's belief and knowledge of it and how these came about.⁴⁶

For reflective knowledge you need moreover an epistemic perspective that licenses your belief by its source in some virtue or faculty of your own. You trust your own correctness, holding your belief to be right through its origin in a reliable faculty or virtue.⁴⁷

For reflective knowledge one not only must believe out of virtue. One must also be aware of doing so.⁴⁸

Reflective justification, our best reflective intellectual procedure, is a matter of perspectival coherence – and necessarily so.⁴⁹

... “Reflective” knowledge [requires] not only apt belief but *also* defensibly apt belief.⁵⁰

Reflective knowledge goes beyond animal knowledge, and requires also an apt apprehension that the object-level perceptual belief is apt. What competence might a believer exercise in gaining such meta-apprehension? It would have to be a competence enabling him to size up the appropriateness of the conditions.⁵¹

... Reflective knowledge requires a specific *further* condition, namely perspectival endorsement of the reliability of one's sources.⁵²

Attaining [reflective] knowledge requires a view of ourselves – of our beliefs, our faculties, and our situation – in the light of which we can see the sources of our beliefs as reliable enough.⁵³

⁴⁶ Sosa 2001: 240.

⁴⁷ Sosa 2001: 277.

⁴⁸ Sosa 2001: 278.

⁴⁹ Sosa 2001: 291.

⁵⁰ Sosa 2007: 24.

⁵¹ Sosa 2007: 108.

⁵² Sosa 2009: 136.

Reflective knowledge is animal belief aptly endorsed by the subject. We can now see that knowing something full well requires that one have animal and reflective knowledge of it, but also that one know it with full aptness.⁵⁴

[Reflective] justification is acquired through rational endorsement, at least in part. It requires the rational endorsement of the reliability of the competence exercised, or of the appropriateness of the conditions for its exercise, or both.⁵⁵

As each of the preceding passages seems to illustrate, Sosa intends for reflective knowledge be a different kind of knowledge than *mere* animal knowledge. Sosa's epistemically undemanding conception of animal knowledge offered the resources needed to make sense of many of the beliefs we commonsensically think we know. However, the epistemically demanding account of reflective knowledge is aimed at a higher goal: an account of knowledge that is philosophically satisfying.

Because Sosa seems to want to make reflective knowledge more philosophically satisfying than mere animal knowledge, one might think that he would be opposed to characterizing reflective knowledge as just *more* animal knowledge. However, Sosa's attempts to further articulate his account of reflective knowledge illustrate some tension. Sosa's presentation of reflective knowledge varies between being centered upon a notion of coherence (or the relevant perspectival endorsement), and being centered upon the notion of moving up an epistemic level (apt belief that is aptly believed to be such).

⁵³ Sosa 2009: 147.

⁵⁴ Sosa 2011: 11.

⁵⁵ Sosa 2011: 150. This passage seems especially telling, given the tension I am trying to develop in Sosa's account. Sosa seems most interested in developing an account of reflective knowledge that requires both rational endorsement, and the appropriateness of conditions. But as this passage makes clear, he is hesitant to make such a commitment.

While the former characterization adds a further condition of coherence onto the requirements set forth for animal knowledge, the latter characterization does not require anything different from just more animal knowledge:

If K represents animal knowledge and K^+ reflective knowledge, then the basic idea may be represented thus: K^+p iff KKp .⁵⁶

2.1.2.1 The Problem of Easy Knowledge

The internal tension that results from the competing conceptions of reflective knowledge that Sosa offers may be the result of another problem that has occupied Sosa's attention; the problem of easy knowledge. As Richard Fumerton and Stewart Cohen have famously pointed out, the reliabilist seems to be able to dismiss skeptical worries too easily. Consider Cohen's illustration:

Suppose I have reliable color vision. Then I can come to know, e.g., that the table is red, even though I do not know that my color vision is reliable. But then I can note that my belief that the table is red was produced by my color vision. Combining this knowledge with my knowledge that the table is red, I can infer that in this instance, my color vision worked correctly. By repeating this process enough times, I would seem to be able to amass considerable evidence that my color vision is reliable, enough for me to come to know my color vision is reliable.⁵⁷

This kind of bootstrapping – this kind of picking oneself up by one's laces – is an example of an illegitimate process of reasoning. But it is not clear why, according to the reliabilist, such a process is illegitimate. After all, the initial beliefs were formed reliably (by hypothesis), and if repeated, the inductive process would also allow us to conclude the general truth that our color vision is reliable. If one were worried about the reliability of one's color vision, would such an argument be satisfying?

⁵⁶ Sosa 2007: 32.

⁵⁷ Cohen 2002: 316.

Sosa's account of robust reliabilism may initially seem especially well-suited to answer the critic who is worried about the problem of easy knowledge. After all, couldn't Sosa maintain that while it is true that animal knowledge may come very easily, the goal of reflective knowledge is to be a kind of knowledge not susceptible to the problem of easy knowledge? Surely the intellectually demanding restrictions placed on reflective knowledge – the required epistemic perspective in particular – ensure that reflective knowledge does not come too easily.

In his explication of the nature of justification at the reflective level Sosa articulates the kind of perspective required for reflective knowledge:

Reflective justification is web-like, not transmissively linear. The web of belief attaches to the world through perception and memory. But each of those nodes depends on other nodes directly or indirectly. The web is woven through the rational basing of beliefs on other beliefs or experiences. There is no reason why such basing must be asymmetrical, however, no reason that precludes *each* belief from being based at least in part (perhaps miniscule part) on other beliefs. *Each* might thus derive its proper epistemic status from being based on others in a web that is attached to the world by causation through perception or memory.⁵⁸ Epistemic justification works more like a web than like a pipe that transmits the juice of justification or warrant...Through the basing of beliefs on other beliefs and on experiences, a rational web is woven, each member of which is upheld in part (perhaps in miniscule part) by others, directly or indirectly. ... Through our growing knowledge of ourselves and of the world around us and of the relation between the two, we come to see our modes of rational basing and other belief acquisition as sufficiently reliable.⁵⁹

In Sosa's presentation of reflective knowledge one can see his attempt at capturing the internalist intuition that there must be more to knowledge than what a simple kind of externalism has to offer. Thus, in response to the problem of easy

⁵⁸ Sosa 2009: 22.

⁵⁹ Sosa 2011: 150–1.

knowledge, Sosa argues that we are able to eventually reach (and enhance) reflective knowledge, “as we gain the ability to explain in richer and richer explicit detail just how it all works to our epistemic advantage.”⁶⁰

One can see the importance of the requirement of perspectival coherence and the role it plays for reflective knowledge. However, it is at this point one might wonder what to do with Sosa’s competing characterization of reflective knowledge. What about reflective knowledge that comes by way of apt belief that is aptly believed to be such? We can, by way of perception, for instance, have animal knowledge about the world around us. But then, on Sosa’s view, it is possible that we have animal knowledge that our perceptual faculties are reliable (recall the K^+p iff KKp principle). If we have animal knowledge that the deliverances of our perceptual faculties are reliable, then it seems we have too easily constructed the perspective needed for reflective justification.

On the apt-belief-aptly-believed-to-be-such characterization of reflective knowledge, we can see once again the threat of the problem of easy knowledge. The intellectual demanding and satisfying account of reflective knowledge gives way to an account that is more attainable, to be sure, but as the problem of easy knowledge illustrates, reflective knowledge, like animal knowledge, comes too easily.

Sosa is not oblivious to the problem of the easy knowledge, nor does he ignore it. We will return to this issue, and to an evaluation of the efficacy of Sosa’s response in the Chapter 4.

⁶⁰ Sosa 2004: 305.

2.1.2.2 Results from the Demanding Level

Setting the preceding discussion aside, we can return to an explication of Sosa's ultimate aim in developing an account of reflective knowledge. Sosa's attempt at achieving epistemic ascent can be seen as an exploration of how an essentially externalist account of knowledge can satisfy the internalist interlocutor who continually asks, "But how do I *know*?" Sosa's response is that, from the animal level, one is not in the right position to answer the question. However, if one ascends to the higher reflective level, then one is in a position to do so:

At an unreflective level, epistemic justification can hence derive from the holding of a condition whose absence is no more subjectively distinguishable from its presence than is a realistic dream from waking life. Still, without reflective, non-arbitrary assurance that you satisfy that condition, you cannot know reflectively something you might still know at the animal level.⁶¹

2.2 Sosa's Theory Applied

2.2.1 The New Evil-Demon Problem

Consider again the new evil demon problem as it was introduced in Chapter 1. Imagine two different cases. In the first case, "unknown to us, our cognitive processes, those involved in perception, memory and inference, are rendered unreliable by the actions of a powerful demon or malevolent scientist."⁶² In the second case, the very same cognitive processes are not affected by a powerful demon; instead our processes work just as they should, reliably, in our environment. By hypothesis in the two cases the background beliefs are the same, and the sensory experiences are phenomenally indistinguishable from one another.

⁶¹ Sosa 2007: 16.

⁶² Lehrer and Cohen 1983: 192–3.

As we have already seen, the new evil demon problem can play a prominent role in motivating certain varieties of internalism.⁶³ Furthermore, the point of the new evil demon problem is not a skeptical one. The problem seems to be one of trying to accommodate the seemingly uncontroversial assumption that the justificatory status of the beliefs of both individuals should be the same. Sosa develops the new evil demon problem in the following way:

The victim of Descartes' evil demon seems *not* deprived of *ordinary* justification, in some straightforward sense, since his beliefs still derive from sources that we recognize as justification-conferring: namely, sense experience, memory, et cetera.⁶⁴

As Sosa recognizes, if a person in a non-demon world has a certain set of justified beliefs, then that person's psychological twin in the demon world should have the same set of justified beliefs. The challenge for the externalist, who says the justificatory status of belief is dependent on factors external to the believer, is to accommodate the uniformity of justification across possible worlds (i.e., to explain how the justificatory status of two believers is the same when one is in a demon world and the other is in a non-demon world).

A wide array of responses to the problem have been attempted on behalf of externalism. In what follows I will briefly sketch some alternative approaches to answering the new evil demon problem and then discuss Sosa's preferred response to the problem. Broadly speaking, responses to the new evil demon problem can be divided into two groups: those that deny the intuition that underlies the problem, and those that try to reconcile the intuition with reliabilism.

⁶³ See Chapter 1, section 1.

⁶⁴ Sosa 2009: 35.

2.2.1.1 Denying the Underlying Assumption

The intuition in question (or the “uncontroversial assumption,” as I call it above) is the idea that the justificatory status of the beliefs of the two individuals (the individual in the actual world and his counterpart in the demon world) is the same. If my beliefs in the non-demon world are justified, then the beliefs of my counterpart in the demon world should be justified too.

According to the first kind of response to the new evil demon problem, one need only deny the intuition that the beliefs had by the person in the demon world are justified (while granting that the beliefs had by the person in the non-demon world are justified). In the demon world, it might be argued, the person’s beliefs do not result from a reliable process, so by hypothesis those beliefs are not justified.⁶⁵

While such a response avoids the new evil-demon problem, the cost of denying such a widely held intuition is too great for many. As a result, responses of the second variety, those that try to reconcile the intuition with reliabilism, are the most common.

2.2.1.2 Strong and Weak Justification

Goldman has argued that the relevant intuition can be retained if one accepts a distinction between two kinds of justification: strong justification and weak justification.⁶⁶ According to Goldman, strong justification is what results from a reliable process working in an appropriate environment: a belief that is well formed. Weak justification amounts to belief that one is faultless, blameless, or that an agent would be non-culpable for holding (independent of any requirement of proper formation).

⁶⁵ See Sutton 2007.

⁶⁶ Goldman 1988.

When Goldman applies his distinction between strong and weak justification to the new evil demon problem, he contends that “the victim of the demon fails to have strongly justified beliefs, but he does have weakly justified beliefs. While his beliefs are not well formed, they are blameless and non-culpable.”⁶⁷ The foregoing analysis, according to Goldman, allows us to make sense of the “strong temptation to say that a cognizer in a demon world *does* have justified perceptual beliefs.”⁶⁸

But despite Goldman’s claims to the contrary, it is not so clear that this approach does accommodate the intuition that the subjects in the two worlds both have beliefs with the *same* justificatory status. Goldman has called both concepts “justification,” but if reliability is the essential feature of strong justification, it is at least a stretch to call something else “justification” that is wholly lacking that essential feature.

2.2.1.3 Normal-World Reliabilism

An alternative approach to accommodating the relevant intuition, defended for a time by Goldman, is normal-world reliabilism.⁶⁹ According to normal-world reliabilism, a belief is justified if and only if it was produced by a reliable process in a normal world. Clearly, the crucial notion here is what counts as a normal world. According to Goldman, it is a world that is, very roughly, the same as (or very similar to) the way we think the actual world is. Thus, normal worlds “are worlds consistent with our *general* beliefs

⁶⁷ Goldman 1988: 60.

⁶⁸ Goldman 1988: 59.

⁶⁹ Goldman 1986.

about the actual world.”⁷⁰ Goldman contends that our “general beliefs” about the actual world are “beliefs about the kinds of things that, realistically, do and can happen.”⁷¹

Given the important role that our general beliefs play in constraining which reliable processes lead to justified beliefs, one might hope that such general beliefs are at least well founded. However, as others have pointed out, the normal-worlds proposal does not constrain the formation of these general beliefs, and as a result, one can imagine different sets of general beliefs that would lead to wildly different sets of justified beliefs.⁷²

2.2.1.4 Transglobal Reliabilism

Henderson and Horgan have argued that there is another way of making sense of the intuition that motivates the new evil demon.⁷³ They argue that a belief is justified iff it is a product of processes that are transglobally reliable. They define processes as transglobally reliable when “the belief-fixing process [is] reliable in a wide range of such global environments.”

Henderson and Horgan use a distinction between local and global reliability to help illustrate their crucial concept of transglobal reliabilism. Imagine asking the reliabilist, *Are one’s perceptual faculties reliable in Fake-Barn County?* At the local level (the level which takes into account one’s immediate environment) it would seem the reliabilist is committed to saying that one’s perceptual faculties are not reliable. However,

⁷⁰ Goldman 1986: 107.

⁷¹ Goldman 1986: 107.

⁷² Pollock and Cruz 1999.

⁷³ Henderson and Horgan 2007: 102.

if we broaden the environment to the global level (the level which takes into account one's entire environment), we can see why the reliabilist might conclude that one's faculties are reliable. Henderson and Horgan want to broaden the environment that determines reliability even further: "... The relevant form of reliability is reliability relative to the set of *experientially possible global environment*..."⁷⁴ Thus, when considering the status of one's perceptual faculties in Fake-Barn County, it is neither the local level nor the global level that determines reliability – it is the transglobal level.

We can apply the resources of transglobal reliabilism to the new evil demon problem. According to transglobal reliabilism, the fact that an agent might be in an evil-demon world does not affect the justificatory status of her beliefs. Why? Because reliability is determined transglobally (i.e., not *only* globally: according to the world that we happen to be in, but transglobally: according to the set of experientially possible global environments). The justificatory status of a person's beliefs in the actual world would be the same as in the evil-demon world (on the assumption that the two people have identical beliefs produced by the same process). Since the relevant reliability of the process that determines justification is transglobal reliability, it does not matter for justification which particular world a person is in.

Transglobal reliabilism certainly offers a novel approach to solving the new evil demon problem. However, I do not find the essential feature of this account plausible, because it is unclear to me how to understand the idea of a belief-producing process that is reliable in *most* of the experientially possible global environments. I am suspicious of any such endeavor that might result in the conclusion that any of our belief-producing

⁷⁴ Henderson and Horgan 2007: 101.

processes are reliable transglobally. For each experience that I have in the actual world, it would seem there is an infinite number of worlds that are experientially possible. I am simply not convinced that any belief-producing process in the actual world is reliable in *most* of those experientially possible global environments. In fact, it seems to me, that there are many more experientially possible global environments where things go wrong, as opposed to the few experientially possible global environments where things go right.

2.2.1.5 Actual-World Reliabilism

Sosa's response to the new evil demon problem also depends on the nature of the world in which the relevant belief is being formed. The view that Sosa first developed, and Comesaña further articulated,⁷⁵ depends on the idea that one must carefully distinguish between two kinds of reliabilism:

Generic Reliabilism: the belief that P is justified iff P was produced by a process that is reliable.

Actual-World Reliabilism: the belief that P is justified iff P was produced by a process that is *actually* reliable.⁷⁶

Sosa and Comesaña make explicit in their characterization of actual-world reliabilism that our claims about the justification of particular beliefs are themselves relative to an environment. Thus, according to the reliabilist, “‘actual’ itself is an indexical term: for any world *w*, an utterance of ‘actual’ in *w* refers to *w*.”⁷⁷

Taking justification to be relative to a certain environment, Sosa contends: Relative to our actual environment A, our automatic experience-belief mechanisms count as virtues that yield much truth and justification. Of

⁷⁵ I take Sosa's view to be essentially the same as that of Comesaña 2002.

⁷⁶ This characterization of actual-world reliabilism is based on what Comesaña calls *indexical reliabilism*. See his 2002: 256.

⁷⁷ Comesaña 2002: 256.

course, relative to the demonic environment D such mechanisms are not virtuous and yield neither truth nor justification. It follows that relative to D the demon's victims are not justified, and yet *relative to A their beliefs are justified*. Thus may we fit our surface intuitions about such victims: that they lack knowledge but not justification.⁷⁸

According to Sosa, the internalist intuition that the justificatory status of the subject's beliefs in the demon world is the same as that of the subject's beliefs in the non-demon world is a result of judging the subject in the demon world from the perspective of the non-demon world. Sosa would argue that the beliefs of the victim of the Cartesian demon are not justified relative to *her* environment, but relative to *our* environment, they are.

The present proposal, like normal-world reliabilism, assumes that beliefs in the actual environment are reliable. In effect, Sosa's response depends on the idea that our automatic experience-belief mechanisms are reliable in our environment. However, as others have pointed out, this is worrisome:

... Suppose we are victims of an evil deceiver, or that we are brains in a community vat. Then [the mechanisms of another victim in a demon world] are no more reliable in our environment than they are in hers. And thus, according to Sosa's account, [her] beliefs are not justified relative to our environment.⁷⁹

Again, to his credit, Sosa is well aware of the challenges to his response to the new evil demon problem. We will return to this issue, and to the efficacy of Sosa's response, in Chapter 4.

2.2.2 Knowledge of the External World

On Sosa's view, it is easy enough to get animal knowledge of the external world. In order for me to have animal knowledge that there is a cup of coffee in front of me, it

⁷⁸ Sosa 2001: 144.

⁷⁹ Greco 1993: 419.

must be the case that my belief about the coffee is true, and that the accuracy of my belief be the result of a reliable process – in this case, perception. But, one might wonder, how can we have reflective knowledge of the external world? How might one come to know that one’s belief-forming processes are reliable?

Put more succinctly, according to Sosa’s externalism, *if* my perceptual faculties are reliable, *then* my perceptual beliefs constitute knowledge.⁸⁰ But, one might wonder, *How do I know that my perceptual faculties are reliable?* This question is not a request for more animal knowledge, but a request for a description of what is required for reflective knowledge of the external world.

According to Sosa, one gains the appropriate epistemic perspective to license reflective knowledge of the external world by recognizing the internal coherence of her sensory experiences, or her beliefs about such experiences (in conjunction with the rest of her beliefs).⁸¹ Of course, at this point, one might worry along with Sosa that mere internal coherence is insufficient for knowledge. Imagine the case of a crystal-gazer who has perfectly internally coherent beliefs. Why, we might ask, doesn’t the crystal-gazer have knowledge?

Sosa argues that while the crystal-gazer may attain coherence and justification, she does not gain knowledge:

On this view, the crystal-gazers differ from the perceivers in that gazing is not reliable while perceiving is. ... The perceivers *can* know their theory

⁸⁰ At least, in the cases where the further conditions are met that the belief is true and that it is true because of the epistemic virtue in question.

⁸¹ Thus, for example, we need not worry about dreaming because, “in a dream there would be signs to the contrary (recall Austin and Descartes)...” Sosa 2007: 111.

to be right when they know it in large part through perception, since their theory *is* right and perception *can* thus serve as a source of knowledge.⁸²

Sosa will argue that not only do we know that *if* our perceptual faculties are reliable *then* our perceptual beliefs constitute knowledge (in virtue of his externalism), but we also know *that* our perceptual faculties are reliable. We will have the proper perspectival coherence in many cases of perceptual belief, and we are also *prima facie* justified in the belief that our senses are reliable:

Epistemically justified trust in our sensory sources is a gift of natural evolution, which provides us with perceptual modules that encapsulate sensory content and reliability in a single package.⁸³

We accept [the senses'] deliverances at face value as a default stance, properly so.⁸⁴

...Our senses enjoy a kind of default rational justification denied to (ordinary) instruments. That is to say, we are justified in accepting the deliverances of our senses, but we need a rational basis for accepting the deliverances of our instruments.⁸⁵

As a result of Sosa's positive view regarding the *prima facie* justification-conferring state of our senses, he is able to shrug off the criticism leveled against other externalists: Sosa goes beyond animal knowledge of the external world and argues that we can rationally know things about the world around us as well.

⁸² Sosa 2009: 200.

⁸³ Sosa 2011: 137–8.

⁸⁴ Sosa 2011: 138.

⁸⁵ Sosa 2011: 139.

CHAPTER 3

PRESERVING INTERNALISM

3.1 Motivating Internalism

In light of the prominent role that the internalism/externalism debate has played in contemporary epistemology, it should not be surprising that one would find different attempts at motivating internalism in the philosophical literature. Additionally, the different ways of carving up the internalism/externalism controversies naturally lend themselves to different arguments in favor of different kinds of internalism. Adequate motivation of access internalism, for example, might not adequately motivate internal-state internalism. As a result, the evaluation of an argument in favor of internalism will depend on the variety of internalism that one is trying to motivate.

The project of this chapter is to first explicate the most popular strategies for motivating internalism. Second, I will argue that *the* critical element in motivating the most interesting kind of internalism relies on the significance of doing philosophy from the first-person perspective. Third, I will argue that it is from the first-person perspective that internalism is best understood, and as a result, the first-person perspective is crucial to adequately appreciating the kind of internalism that I will defend: assurance internalism. Finally, I will develop an account of noninferential justification that is consistent with assurance internalism.

3.1.1 The Normative Conception of Justified Belief

One common type of argument for internalism relies on a normative (or, following Goldman's terminology, a guidance-deontological) conception of justification. On this view, it is a subject's (epistemic) duty to be guided by his evidence (i.e., the

guidance constraint) and justified belief amounts to belief that is permissible given the subject's evidence, while unjustified belief is belief that is impermissible (or forbidden) given the subject's evidence (i.e., the deontological constraint).

Once this kind of guidance-deontological conception of justification is posited, the internalist might argue that if justification is a matter of praising or blaming a subject (or a subject's beliefs) then it would only seem fair to make this judgment according to the evidence that the subject has access to. On this view, one might try to motivate internalism in the following way:

- P1. The *guidance-deontological conception of justification* is posited.
- P2. A certain constraint on determiners of justification is derived from the guidance-deontological conception, that is, the constraint that all justification determiners must be *accessible to*, or *knowable by*, the epistemic agent.
- C. The accessibility or knowability constraint is taken to imply that only internal conditions qualify as legitimate determiners of justification. So justification must be a purely internal affair.⁸⁶

Goldman is not alone in taking the primary motivation of internalism to be a commitment to normative conceptions of justification, nor is he the only philosopher who endorses a normative conception of justification. Normative conceptions of justification have been advanced by internalists and externalists alike:

...being justified in believing that p consists in some sort of 'deontological' status, for example, being free from blame for believing that p or having satisfied one's intellectual obligations in doing so.⁸⁷

...epistemic justification is essentially a matter of duty fulfillment.⁸⁸

⁸⁶ Goldman 1999 advances this argument on behalf of the internalist, but of course, does not endorse the argument.

⁸⁷ Alston 1991: 72–3.

⁸⁸ Bergmann 2000: 87.

...the concept of epistemic justification is fundamentally a normative concept. It has to do with what one has a duty or obligation to do, from an epistemic or intellectual standpoint.⁸⁹

...epistemic deontology [is] the view that epistemic duty and obligation are of crucial epistemic importance and that ... *being* [epistemically] *justified* is being within our rights, flouting no epistemic duties, doing no more than what is permitted ... [and being] subject to no blame or disapprobation.⁹⁰

Epistemic deontology is the view that the concept of epistemic justification is deontological: a justified belief is, by definition, an epistemically permissible belief...⁹¹

If one held a normative conception of justification, then it is at least plausible why one might further suppose that the determiners of justification must be limited to what is, in some sense, available to the subject.⁹² Intuitively, if one holds that some condition is relevant to justification, but it is impossible to know whether that condition obtains (or is likely to obtain), then, one can hardly be blamed if that condition is not met.⁹³ Of course, the success of this kind of argument for internalism depends on the viability of normative conceptions of justification.

3.1.1.1 Rejecting the Normative Conception

It is at least obvious that we *talk* about justification as if it involves a certain kind of normativity. Epistemologists often switch between talking about what beliefs are

⁸⁹ Bonjour 1980: 55.

⁹⁰ Plantinga 1993: 13–14.

⁹¹ Steup 2000: 25.

⁹² For a detailed defense of why normative conceptions of justification require some sense of ‘availability,’ see Bergmann 2006: 89–98.

⁹³ It is at least worth noting that the intuition behind this kind of argument seems to fall in line with another issue in contemporary moral theory: am I responsible for actions that involve moral luck?

justified and what beliefs *should* be held. However, the fact that people (including epistemologists) talk this way does not mean that on the final analysis we should *understand* justification this way.

A successful defense of the claim that justification is essentially a normative concept will require much clarification. If normative concepts in ethics are analogous to normative concepts in epistemology, then the advocate of normative conceptions of justification would seem to face all of the problems that occupy a central place in contemporary metaethics. Should normative claims be understood as descriptive claims? Or, should normative claims be understood as prescriptive claims? If normative claims are taken to be descriptive, are the truths they express objective or subjective?

The foregoing comments are not meant to imply that these questions cannot be answered. Instead, it is only meant to illustrate the challenges that face a normative conception of justification. Perhaps though, these challenges should make one reconsider whether, simply because we talk this way, the normative conception of justification ought to be pursued.⁹⁴

If the kind of argument Goldman presents on behalf of the internalist is sound, then normative conceptions of justification seems to entail a certain kind of internalism. However, there is what I take to be a fundamental problem with normative conceptions of justification. There seem to be two distinct questions someone might ask. The first question asks whether a particular belief is justified. The second question asks whether the believer is praiseworthy (or blameworthy) in holding a given belief. Since these two questions are different, it seems to me that whether we should praise (or blame) a believer

⁹⁴ Of course, the normative conception of justification might be the result of other considerations too.

does not seem to does not decide the question of whether the believer is justified (or unjustified). Praise (or blame) is secondary to, not constitutive of, justification. First, one must decide if a belief is justified, and in light of this, the person can decide whether to praise or blame the believer. If this is right, then epistemic evaluation is different from and more fundamental than normative evaluation.⁹⁵

To be clear, my worry is that a person is praiseworthy (or blameworthy) because her belief is justified (or unjustified). Being praiseworthy (or blameworthy) is dependent on the question of whether the belief is justified. Put another way, the normative evaluation is separate from and dependent on the concept of justification.

If this line of reasoning is correct and the normative conception of justification fails to get at the most fundamental concept of justification, then one might worry that the kind of internalism that results from the normative conception of justification would also fail to get at the heart of the internalism/externalism debate. While it would not count as a refutation of normative conceptions of justification, it is at least worth mentioning that many contemporary self-proclaimed internalists and many paradigm historical proponents of internalism do not hold a normative conception of justification.⁹⁶ If the only motivation for internalism is the kind of argument that Goldman considers, then it would seem that many internalists simply hold an unmotivated view.

⁹⁵ On this way of dealing with normative evaluation it is often helpful to distinguish the evaluation of belief from the evaluation of the agent.

⁹⁶ For instance, see the contemporary works of Bonjour, Fumerton, McGrew, and the historical works of Descartes, Hume, and Russell.

3.1.2 Mentalism

Conee and Feldman have presented a number of examples that are intended to motivate a specific version of internalism. In each of their examples Conee and Feldman explicitly specify an *internal* difference that ends up making an *epistemic* difference. For instance, Conee and Feldman ask us to consider the following example:

Bob and Ray have both read in a highly reliable newspaper the forecast for the day's weather. In light of the newspaper's forecast Bob and Ray form the belief that the weather will be very warm today. Bob goes outside and feels the heat, and both Bob and Ray continue to believe that it is very warm today.⁹⁷

Conee and Feldman argue that "Bob's belief is better justified ... the belief was enhanced by his experience of feeling the heat..."⁹⁸ Since the only relevant change in the scenario was Bob's sensation of heat (a new internal state) this example illustrates the general point Conee and Feldman are trying to make: changes in justification are dependent on internal changes in an agent. Each of the cases that Conee and Feldman construct are intended to show that the best explanation of the increase in justification is that the addition of a new mental state makes the justificatory difference between the subjects.

Even if it is true that some new internal state does the best job of the explaining the justificatory difference in Conee and Feldman's examples, it would seem that most (if not all) theories of justification could at least make sense of the justificatory change in the examples. The coherence theorist, who defines justification as coherence amongst beliefs, could argue that the new internal state increased coherence. The paradigm reliabilist, who

⁹⁷ Conee and Feldman 2004.

⁹⁸ Conee and Feldman 2004: 59.

defines justification as belief that results from a reliable cognitive faculty, could argue that the increase in justification was the result of a new internal state caused by a reliable belief-forming process.

If the data of the examples that Conee and Feldman present can be accommodated by any theory of justification, then one might wonder how effective Conee and Feldman's argument is. However, there is another kind of example, as we have already seen, that might help motivate the view that Conee and Feldman wish to defend. Consider again the new evil-demon hypothesis:

Imagine two different cases; in the first case "unknown to us, our cognitive processes, those involved in perception, memory and inference, are rendered unreliable by the actions of a powerful demon or malevolent scientist."⁹⁹ In the second case the very same cognitive processes are not affected by a powerful demon, instead our processes work just as they should, reliably, in our environment. By hypothesis it is stipulated that in the two cases the background beliefs are the same, and the sensory experiences are phenomenally indistinguishable from one another.¹⁰⁰

If one shares the strong intuition that the justificatory status of beliefs about external objects in this world should be the same as the justificatory status of identical beliefs in a world where one is the victim of a Cartesian Demon, then one is in a good position to argue that it is the internal states of a subject that make a justificatory difference. Thus, if the internal states of two subjects are identical, so too will be the justificatory status of their beliefs. It is, *following* Conee and Feldman, the lack of internal differences that explain the lack of an epistemic difference.

I take the new evil-demon hypothesis to offer the best support for a view like Conee and Feldman's. Unlike the Bob and Ray example, the new evil-demon hypothesis

⁹⁹ Lehrer and Cohen 1983: 192–3.

¹⁰⁰ Reconstructed from Lehrer and Cohen 1983.

offers an actual situation that is consistent with a view like Conee and Feldman's but cannot be explained by the kinds of views they are rejecting.¹⁰¹ The reliabilist, for example, seems stuck with the challenge of explaining how reliable beliefs (like ours in the actual world) have the same justificatory status as unreliable beliefs (like those of the victim of the Cartesian demon). Conee and Feldman should argue that while their view can explain the data (that the subject in the actual world and the victim of the Cartesian demon have beliefs with the same justificatory status) the competing views (like reliabilism) cannot, since by hypothesis there is no reliable belief forming process at play for the victim of the Cartesian demon.

Conee and Feldman name their view *mentalism*, and they develop it in the following way:

A person's beliefs are justified only by the things that are internal to the person's mental life... As long as the things that are said to contribute to justification are in the person's mind, the view qualifies as a version of mentalism.¹⁰²

According to Conee and Feldman justification supervenes on the mental states of a believer. Furthermore, Conee and Feldman argue that this supervenience relation holds across possible worlds, and as a result any two subjects in any two worlds who are exactly alike mentally would be exactly alike justificationaly.

3.1.2.1 Rejecting Mentalism

While mentalism might be classified under the label *internalism*, the view fails to capture the most interesting sense of internalism. Mentalism puts the focus of internalism

¹⁰¹ Of course much more needs to be said about whether externalist responses to the new evil-demon hypothesis are adequate. See Chapters 2 and 4 for a discussion of the kinds of moves available to the externalist in response to this kind of problem.

¹⁰² Conee and Feldman 2004: 55.

on the location of justification contributors. By making internalism a matter of where justification contributors are located, Conee and Feldman's version of internalism seems to invite the troubling set of questions that face internal-state internalism. As we have already seen, there are live controversies about whether mental states are, to paraphrase Putnam, *in the head*. The more sympathetic one becomes to semantic externalism the harder it becomes to equate internal states with mental states. Someone who wishes to defend mentalism will have to develop a full-fledged account of internal states.

Conee and Feldman want to stay neutral on the question of what the correct account of mental states is. Their goal is to advance an epistemological thesis: justification supervenes on mental states. They make their claim a purely epistemological one, and refuse to address the metaphysical question of what counts as a mental state. Conee and Feldman invite others to offer the account of mental states that they find most plausible, where the result will be different varieties of mentalism (where each different conception of what counts as a mental state will entail a different kind of mentalism). In order for mentalism to be a substantive and evaluable thesis an account of what a mental state is must be offered. And, as we have seen, opinions as to what counts as a mental state are varied, to say the least.

In answering the question of what counts as a mental state one might be able to come up with a gerrymandered account of what an internal state is. This gerrymandered account will likely include such disparate states as feeling pain, being aware of one's pain, being aware of the surface of physical objects, being acquainted with numbers, and being acquainted with universals. But at this point, when one thinks that all these disparate states are the critical states on which justification supervenes, one might well

wonder whether it isn't something other than the fact that there is some extended concept of 'internal state' that we can invent for them to fall under, or whether we are not missing some other factor that is common to all these different states – some other factor that might be critical to justification.

3.1.3 The Subject's Perspective Objection

Perhaps the most common way of motivating internalism is the presentation of the now well-rehearsed cases that are intended to illustrate the shortcomings of externalism:

Norman the clairvoyant: Norman is a clairvoyant, and he has the ability to reliably form beliefs as a result of this power. Norman has no evidence of any kind for or against the possibility of this kind of power, nor does he have any reason for or against the thesis that he has such a power. One day, Norman comes to believe that the President is in New York City, though he has no evidence for or against this belief. Unbeknownst to Norman, the belief was a result of his reliable clairvoyant power, and the belief is in fact true.¹⁰³

Mr. Truetemp: Suppose a tempucomp, a device capable of generating accurate thoughts about the current temperature, is inserted into Mr. Truetemp's head. Further suppose the tempucomp has been inserted into Mr. Truetemp head without his knowledge. It can be assumed that the tempucomp is reliable, the produced thoughts are true, and Mr. Truetemp has never checked up on the accuracy of his thoughts about the temperature.¹⁰⁴ "All told, this is a very reliable belief-forming process and a properly functioning cognitive faculty."¹⁰⁵

The Serendipitous Brain Lesion: Imagine a subject who suffers from a serious abnormality, a brain lesion. The brain lesion causes all kinds of false beliefs for our subject.¹⁰⁶ "It also causes him to believe, however, that he is suffering from a brain lesion. K has no evidence at all that he is abnormal in this way ... but surely K does *not* know that he is suffering

¹⁰³ Reconstructed from BonJour 1980.

¹⁰⁴ Reconstructed from Lehrer 2000.

¹⁰⁵ Lehrer 2000: 187.

¹⁰⁶ Plantinga 1993: 195.

from a brain lesion... his holding this belief is, from a cognitive point of view, no more than a lucky (or unlucky) accident.”¹⁰⁷

Each of these cases is intended to motivate internalism. However, it is not immediately obvious just how (or even what kind of) internalism is supposed to be motivated by these cases. Bergmann, in an effort to ultimately defeat internalism, has argued that best way to make sense of what is going on in these cases is to emphasize the importance of what he calls the subject’s perspective.¹⁰⁸

According to Bergman, the internalist should point out that in each of the hypothetical cases the subject would have justification according to externalism, but clearly, it is argued, the subjects in each of the cases are not justified. For instance, consider the first case, that of Norman. Norman is stipulated to have a reliably produced true belief in the absence of any defeaters. According to at least one popular version of externalism, reliabilism, Norman’s belief is justified. However, from Norman’s perspective, the belief that the President is in New York City is entirely random.¹⁰⁹ As a result, the externalist seems saddled with having to bite an unsavory bullet, one can have justified belief even when the belief in question seems utterly random.

The foregoing analysis seems to show that the conditions on justification offered by the externalist are too weak – we end up with justified belief too easily. From this realization, Bergmann, on behalf of the internalist, develops this idea in what he calls the Subject’s Perspective Objection (SPO):

¹⁰⁷ Reconstructed from Bergmann 2006.

¹⁰⁸ Bergmann 2006: 11–13.

¹⁰⁹ By saying that the belief appears random to Norman I am trying to point out that, from his perspective, he has no reason for thinking he should have ended up with this particular belief about the president.

SPO: 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 perspective it is an accident that her belief is true. And that implies that it isn't a justified belief.¹¹⁰

Bergmann argues that the primary motivation for (awareness) internalism is the SPO. Furthermore, internalists who advocate some variety of the SPO commit themselves to awareness internalism.¹¹¹ Bergmann proceeds by arguing that awareness internalists, when trying to develop the conception of awareness required on their view, commit themselves to either a vicious regress or a form of awareness that is too weak to capture the SPO. As a result, according to Bergmann, “we should not endorse internalism.”¹¹²

Our task is to come up with philosophically interesting characterizations of both internalism and externalism. They two sides of one coin: if we can get a philosophically satisfying account of internalism, we can define externalism as the rejection of the internalist thesis. Alternatively, we could get a philosophically satisfying account of what makes someone an externalist and then define internalism as the view that rejects the externalist thesis. The account I develop combines both approaches. I will show what internalism requires (and by definition, what externalism rejects), and I will show what externalism does not require (and by definition, what internalism rejects).

¹¹⁰ Bergmann 2006: 12.

¹¹¹ Actually, Bergmann's preferred terminology is 'access internalism,' but my terminology should not make any substantive difference in the debate.

¹¹² Bergmann 2006: 21.

3.2 Assurance Internalism

To his credit, I think Bergmann is right in noticing what internalists take to be so important to an analysis of justification. In order to properly understand the claims of the internalist, one must adopt the first-person perspective. When the subject's perspective is taken, and one considers whether one's beliefs are justified, externalist accounts that allow for justified belief despite any kind of awareness of the justification contributors seem entirely uninteresting.

Of course, from a third-person perspective such externalist accounts may seem to be plausible and interesting ways of evaluating the beliefs of others. But again, the internalist should not be concerned with the third-person perspective when it comes to cashing out what it means to be an internalist. On the view I will be developing, the issue of awareness that Bergmann makes explicit with the SPO should be understood in terms of a sense of assurance that results from justified belief.

According to the kind of internalism that I find most plausible – what I have called assurance internalism – when one has a justified belief, one is in a position to see what the belief has going for it. One will have the resources to satisfy the kind of intellectual curiosity that evades paradigm versions of externalism. As a result, the most helpful way of understanding the essential difference between internalists and externalists depends on this notion of assurance. If one subscribes to an account of justification that does not result in the kind of assurance that, from the subject's perspective, results in the subject's seeing what the belief has going for it, then the account of justification is not an internalist one.

Alternatively, we can define externalism by what it lacks: externalism about justification is the view that one can be justified in a belief P even if one lacks assurance of the truth of P. Or, put another way, an externalist allows for one to have justified belief even if one does not see what the belief in question has going for it.

3.2.1 Motivating Assurance Internalism

Like so many other fundamental philosophical concepts the sense of assurance on which my view depends is exceedingly difficult to characterize. Even though it is difficult to specify the precise nature of an epistemic state like assurance, one can at least use examples to show when a subject obviously has assurance, or when the subject obviously fails to have assurance.

Example 1: Jack, because of his love of mathematical propositions that look complex, adamantly believes that some mathematical proposition that looks really complex is true. Furthermore, we can imagine that the proposition in question does express a necessary truth.

Example 2: Fletcher, who has never followed politics, happens upon a political debate on television. After listening to the first candidate make her case, Fletcher turns off the television and believes all the positions the candidate advocated are true. As it turns out, all the positions advocated by the candidate were true.

Example 3: Lilly believes that God exists, even as she finds herself in the grips of an all-encompassing global skepticism. As it turns out, even though Lilly is in the grips of her skepticism, her belief is true and the result of a reliable faculty functioning in its proper environment.

Example 4: Ryan takes a powerful drug that makes one confident in the truth of any belief one considers. He considers the possibility of extraterrestrial life, and contends that he is certain about his drug-induced belief in extraterrestrials. We can further imagine that Ryan's beliefs turn out to be true.

Example 5: Darren believes that he seems to be drinking coffee, since he is having the experience of drinking coffee.

Example 6: Amy entertains the thought that she is a thinking thing. She believes that since she is a thinking thing, she must exist.

Example 7: McKenna considers a very simple mathematical proposition, and believes it expresses a truth because she can see that it has to be true.

Example 8: Scott believes that all bachelors are unmarried because he understands what it means to be a bachelor.

According to the account of assurance internalism that I am developing examples 1–4 illustrate instances where a subject lacks assurance of their beliefs. Alternatively, examples 5–8 illustrate instances where the relevant sense of assurance is present.

As we can see from the examples, assurance is not necessarily tied to the belief of a necessary truth (example 1). Assurance does not accompany a set of lazily formed beliefs that avoid defeaters (example 2). One might very well lack assurance in the case of reliably formed belief (example 3). Assurance is not simply a matter of psychological certainty (example 4). We can contrast the first six examples with examples 7–10. Examples 7–10 illustrate the built in sense of assurance that accompanies justified belief according to the kind of internalism that I am defending.

3.2.2 Clarifications

It should be made clear that the kind of internalism I am trying to develop puts the central focus of internalism on the sense of assurance that results from justification. However, it might not be clear why the account that I have offered is any different from internal-state internalism (since assurance is certainly an internal state) or awareness internalism.

The account of assurance internalism that I am arguing for differs from internal-state internalism in virtue of the factors that determine whether one is an internalist or externalist. The internal-state internalist, *qua* internal-state internalist, defines internalism by limiting justification contributors to the internal states of the subject. The assurance

internalist, *qua* assurance internalist, defines internalism according to the sense of assurance that results from justified belief. As a result, the views differ in how they make sense of the internalism/externalism debate.

My claim is that by carving up the internalism/externalism debate through appeal to the sense of assurance that is built into a satisfactory account of justified belief, the assurance internalist best reaches the heart of the internalism/externalism controversies.

Imagine the following case:

After watching the *Matrix*, Leah begins to wax philosophic about the possibility of being a brain in a vat where all of her experiences are really the result of something like Descartes's evil demon. Leah worries to herself about whether her belief that she is at a *real* coffee shop, drinking a real cup of coffee, is justified.

On the account I am offering, what makes this kind of case interesting isn't thinking about it from the third-person perspective. Instead, when I consider the case from the first-person perspective, and realize that I could be in Leah's position, I might ask myself, why should I think my beliefs about the external world are justified?

When I find myself in this kind of reflective mental state the pronouncements of externalism seem uninteresting and wholly unsatisfying. Realizing that my beliefs *might* be the result of a reliable belief-forming process does nothing to satisfy my intellectual curiosity. After all, I want to know whether a particular belief is justified – to what extent is that belief firmly grounded? By making internalism a matter of the assurance that accompanies justified belief, the assurance internalist is carving up the internalism/externalism controversy with an eye toward Leah's worry. Perhaps it is this fundamental interest in taking skepticism seriously that motivates assurance internalism.

The kind of question that motivates assurance internalism seems different from the kind of question that the internal-state internalist seems to be asking. The internal-

state internalist may be equally dissatisfied with the externalist's response to the kind of question that Leah asked. However, the source of the dissatisfaction for the internal-state internalist was not the failure to satisfy her philosophical curiosity, but the externalist's claim that Leah's belief might be justified in light of non-mental states. On my view, the internal-state internalist is focusing on the wrong problem *qua* epistemologist. The most interesting question is not *where* (inside the mind, or outside of it) the justification contributor is located (as the internal-state internalist wants us to believe), but the assurance Leah can lack in relation to the justification contributor (according to the externalist).

Making the advantages of assurance internalism over internal-state internalism explicit might lead one to think that assurance internalism is really just a disguised version of awareness internalism. After all, isn't the essential point that the assurance internalist is making really about *awareness* of the justification contributor for a belief? On the view I am advocating, in contrast to awareness internalism, awareness of the justification contributor is not a further condition on justification. Put simply, the awareness internalist seems to be saying that in order for an account of justification to be a form of awareness internalism, that account must involve the condition of awareness of the satisfaction of truth conditions. Thus, if B is the belief and J is the account of justification, then JB is only a version of awareness internalism if the condition A, awareness, is added: JBA.

According to assurance internalism, the sense of assurance that results from justified belief is not a further condition on justified belief. Instead, the sense of assurance will result *whenever* one has a justified belief (on the appropriate analysis of

justification). Because assurance internalists are not committing themselves to a further condition on justification, the kind of problem that Bergmann tries to motivate for internalism does not arise.

Bergmann argues that the awareness internalist, in accepting a further condition on justification, invites a vicious regress: once justification has been modified with the addition of an awareness condition, a belief is only justified if one is aware of the justification contributor. Of course, awareness of a justification contributor seems to involve awareness that one is aware of the justification contributor, and so on *ad infinitum*. However, by explicitly not making the kind of assurance I am interested in a further condition on justification, this kind of regress is avoided.

The assurance internalist is committed to the claim that assurance will result from, and indeed necessarily accompanies, justified belief (when justification is properly understood). In the next section I will sketch an analysis of justification that results in this kind of assurance, and which as a result is properly construed as a version of (assurance) internalism.

The assurance internalist is also committed to the claim that justification and assurance are two separate things. However, one might worry about the possibility of assurance not accompanied by justification. Can't we imagine cases where a subject experiences the kind of assurance that the assurance internalist is interested in without meeting the condition on justified belief?

Imagine Megan who, for no apparent reason, finds herself believing that her daughter is in immediate danger. Megan forms the belief that my child is in danger. By

hypothesis, we stipulate that Megan is in the relevant state of assurance. What is the assurance internalist supposed to say about such a case?

To start, it is not immediately clear what the case is supposed to show. By hypothesis Megan has no justification for her belief, so surely the proposed case does not directly threaten the internalist's conception of justification. Perhaps the worry is something like this: by making internalism a matter of assurance, the internalist has separated internalism from justification. One can have assurance without having justification.

This kind of objection, I think, misunderstands the claim of the assurance internalist. According to the proposed view, an account of justification is internalist if and only if, from the subject's perspective, when the conditions on justification obtain, a sense of assurance results. Thus, discussions of assurance absent justification are at best orthogonal to the matter at hand. Furthermore, the possibility of assurance occurring absent justified belief is at best controversial on the view of justification I will develop.

3.2.3 The Case for Assurance Internalism

In virtue of capturing the idea behind the SPO it is my contention that assurance internalism makes sense of the intuitive appeal of the kind of argument Goldman presents on behalf of the internalist, as well as the cases Conee and Feldman use to motivate mentalism. Finally, assurance internalism adequately situates predominant internalists and externalists in their appropriate places.

Goldman's argument depended on the thesis that justification contributors should be limited to what is accessible to or knowable by the epistemic agent. In making this kind of claim, Goldman rightly characterizes the importance of the subject's perspective.

As a result, Goldman is correct in characterizing the internalist as being primarily concerned with a concept of justification that takes this point of view.

Conee and Feldman develop cases where justified belief is assumed for two subjects. One of the subjects gains a new mental state that results in a change in their degree of justification (or justificatory status). Conee and Feldman contend that since a change in the relevant mental state resulted in a change in justificatory status, we should conclude that justification supervenes on the mental. However, the very same cases can be used again to illustrate the importance of the subject's perspective. Since the addition of a relevant mental state changed the subject's perspective, it is not surprising that a new justificatory state would result as well.

Finally, in developing a new way of conceiving of the differences between internalism and externalism, it would at least seem odd if paradigm advocates of internalism or externalism failed to fall within their respective categories. Fortunately, the account of the assurance internalism that I have developed does not fall prey to this problem. According to assurance internalism, it would be natural to place Descartes, Hume, Russell, and Chisholm within the category of internalists. Each was interested in a concept of epistemic justification that had something to do with the subject's "seeing" what the belief had going for it. Alternatively, Goldman, Nozick, and Plantinga would all fall under the externalist label. Their conception of justification allows for justified belief in the absence of "seeing" what the belief has going for it.

3.3 Acquaintance and Noninferential Justification¹¹³

I have argued that an account of justification is internalist if, when the conditions on justification are met, a sense of assurance with regard to the belief in question results. In what follows I will develop a recursive analysis of justification that satisfies these criteria. Like any complete recursive analysis of justification, both a base clause and a recursive clause must be developed. For now, I will focus on developing the base clause of the recursive analysis: an account of noninferential justification.¹¹⁴

The account of noninferential justification that I find most plausible depends on two controversial components. The first is the *correspondence theory of truth*. On my conception of the correspondence theory, the primary bearer of truth is thought: a non-relational property of the mind capable of being either true or false. Put another way, my belief that I am sitting in front of a fireplace is a kind of thought that the world is a certain way. The truth maker of a thought is a fact: a non-linguistic complex that exemplifies a set of properties. A fact, then, is a way the world is. For instance, it could be the case that I am sitting in front of a fireplace.

To say that a thought is true is to say that the thought corresponds to (pictures, matches up with, fits) a fact. Alternatively, false thoughts fail to stand in the relation of correspondence to a fact. Thus, my belief that I am sitting in front of a fireplace is false

¹¹³ The role of acquaintance in justification is not limited to the kind of account I develop below. For other accounts of justification that depend on acquaintance, see Russell 1912, Lewis 1929, Price 1950, Fales 1996, and BonJour (BonJour and Sosa 2003). The account of noninferential justification that I develop is the view originally articulated by Richard Fumerton. Fumerton makes the case for his account in numerous places; the account I offer will depend heavily on his 1995 and 2006.

¹¹⁴ Chapter 4 will explore how an assurance internalist might develop a recursive clause of justification that is consistent with the base clause developed here.

because the relation of correspondence does not hold between the belief (a certain kind of thought) and a fact (the way that the world is).

The second relation I wish to defend is the relation of direct acquaintance. Acquaintance, following Fumerton, is “a *sui generis* relation that holds between a self and a thing, property, or fact.”¹¹⁵ To stand in the relation of acquaintance with a thing is to have that thing directly before one’s consciousness. In other words, being acquainted with a thing is to stand in an unmediated relation to the thing.

On this view acquaintance is not, like thought, an intentional state. As a result, the relation of acquaintance alone does not entail noninferential justification. Instead, one is noninferentially justified in one’s belief that P when:

1. One has the thought that P,
2. One is acquainted with the fact that P, and
3. One is acquainted with the relation of correspondence holding between the thought that P and the fact that P.

For example, I have the belief that I seem to be sitting at my desk. In virtue of having the belief that I seem to be sitting at my desk I have the *thought* that I seem to be sitting at the desk (belief being a species of thought). I am acquainted with my *seeming* to be sitting at my desk, and I am acquainted with the relation of correspondence holding between my thought and my seeming. Since I meet all three conditions on noninferential justification, my belief that I seem to be sitting at my desk is justified.

¹¹⁵ Fumerton 1995: 74.

3.4 Acquaintance, Noninferential Justification, and Assurance

I have argued that the most philosophically interesting conception of internalism is assurance internalism (section 2). I have also presented the account of noninferential justification that I find most plausible (section 3). In what follows I argue that the acquaintance theory of noninferential justification satisfies the conditions of assurance internalism, and as a result is best understood as a variety of internalism about justification.¹¹⁶

Assurance internalism was defined as the view that when one has a justified belief one will have a sense of assurance of the truth of that belief. In order to see why the acquaintance theory of justification counts as a variety of assurance internalism, imagine an instance where the conditions on justification (according to the acquaintance theorist) are met. My belief that I have a headache is justified when:

1. I have the thought that I have a headache,
2. I am acquainted with the fact that I have a headache, and
3. I am acquainted with the relation of correspondence holding between my
thought and the fact.

In this example, if my belief is justified, then by definition I meet conditions 1–3. In virtue of being acquainted with the relation of correspondence holding between my thought and fact it follows that I can see what my thought (the belief) has going for it. I have before my consciousness the truth bearer (the thought) and the truth maker (the

¹¹⁶ It is worth noting that I am not arguing that the acquaintance theory of justification is the only possible theory of justification that would count as a version of assurance internalism.

fact), and I am acquainted with the relation of fitting obtaining (the correspondence) between the truth bearer and the truth maker.

On my view, it is in virtue of being acquainted with the relation of fit obtaining between the thought and the fact that results in the sense of assurance that is of interest to the internalist. As a result of seeing that one's evidence (the fact) matches up to (corresponds to) one's belief (the thought), one cannot help but see what the belief has going for it.

If one did not see what the belief had going for it, one would not, by definition, be acquainted with a necessary part of justification. Potentially, one could fail to meet the standards of justification in three different ways while having a thought: in the case where condition 1 is met, a belief is still unjustified if condition 2 fails to be satisfied, condition 3 fails to be satisfied, or conditions 2 and 3 both fail to be satisfied.

In the case where condition 3 is satisfied and I am able to "see" the relation of fitting between my thought and the fact, it is my contention that a sense of assurance necessarily results.

3.4.1 Objections and Replies

To say that the account I have defended here is controversial would be an understatement. Specifically, the conceptions of acquaintance and correspondence on which my view depends are wildly unpopular. In addition, these views commit me to other unpopular theories, for example, within the philosophy of mind. However, as I have argued, the view I defend does capture what I take to be a common historical thread amongst what epistemologists have been looking for: a philosophically satisfying account of justification.

Despite the unapologetic use of controversial theses in my defense of assurance internalism, there is at least one kind of objection that has been aimed at the acquaintance theory of noninferential justification that I have defended myself against. If this objection were successful, it would undercut my use of acquaintance to capture the kind of assurance internalism that I am interested in. This objection comes from Bergmann's explicit challenge to the kind of claim I have been making: that the acquaintance theory of noninferential justification has built into it a sense of assurance.

For the sake of argument, Bergmann grants that one can satisfy each of the three criteria required on the acquaintance account I defend: one's belief that P is noninferentially justified if and only if one is acquainted with one's thought that P, the fact that P, and the relation of correspondence holding between the thought that P and the fact that P. Bergmann asks us to imagine the following example of Jack:

Suppose Jack is being appeared to redly and that he believes that he is being appeared to redly (call this belief 'B2'). ...He can be directly acquainted with the relation of correspondence holding between his thought that he is being appeared to redly and the fact that he is being appeared to redly *even if he has no idea* that the relation of correspondence holds between these two items. Thus, Jack's belief that B2 can satisfy [the acquaintance theorist's requirements] even if he conceives of his being appeared to redly as no more relevant to B2 than is the mild pain in his left knee.¹¹⁷

Because of the possibility of this kind of example, Bergmann concludes that "It is, therefore, exceedingly difficult to see how these direct acquaintances improve things *from Jack's subjective perspective*."¹¹⁸ Bergmann contends that the admittedly odd example of Jack's being acquainted with the correspondence between his thought that P

¹¹⁷ Bergmann 2006: 28–30.

¹¹⁸ Bergmann 2006: 30.

and the fact that P, without having any idea that the fact that P is relevant to his thought that P, is hard to imagine. On my view, this kind of example is hard to imagine for a good reason: it is not possible.

The beauty of thought experiments is that an author can stipulate the conditions of a thought experiment however he or she wants. Bergmann takes this license when he stipulates the conditions that surround Jack. We are told to imagine a case where each of the criteria required for noninferential justification on the acquaintance theory are met, and then told to imagine of the case that it is also true that Jack conceives of his being appeared to redly as no more relevant to his belief that he is being appeared to redly than is his mild pain in his left knee.

If an acquaintance theorist were inclined to think that such an example was possible then the acquaintance theorist would be in trouble. Of course, if the hypothetical scenario Bergmann discusses is incoherent, then the acquaintance theorist has nothing to worry about (at least in regards to the supposed problem from Bergmann). It is my contention that Bergmann's supposedly problematic hypothetical case fails to be problematic because the kind of situation Bergmann describes cannot obtain (at least not when one properly understands the relevant notions of acquaintance and coherence).

A case in which I have (1) a truth maker directly before my consciousness, (2) a truth bearer directly before my consciousness, and (3) I am able to see the fit of the truth bearer to the truth maker (i.e., I have the relation of correspondence directly before my consciousness), and *yet* I conceive of the truth maker as being irrelevant to the truth bearer strikes me as utterly incoherent. How could I, after all, be acquainted with the

relation of fit between the truth maker and truth bearer unless I was actually acquainted with the *fit*?

On the kind of view I am defending, being acquainted with the relation of correspondence holding between a thought and a fact requires that one *sees* that a certain kind of relation *obtains*: the subject can see that the fact corresponds to the thought. As a result, when a subject meets the criteria required of noninferential justification by the acquaintance theory, the subject cannot help but *see* what her belief has going for it. The subject with a noninferentially justified belief (on the acquaintance theory) has everything she could want from the first-person perspective. Thus, in contrast to the example that Bergmann asks us to imagine, one cannot be acquainted with the relation of correspondence *holding* between two things, without *also seeing* the relevance of the one thing to the other.

Bergmann could argue that the conceptions of acquaintance and correspondence on which this view depends are implausible, but of course, that would be a different kind of argument.

CHAPTER 4

ACHIEVING EPISTEMIC DESCENT

4.1 Rejecting Sosa's Account

The general architecture of Sosa's epistemology is designed to distinguish between kinds of knowledge. While I take his general strategy of making this distinction (or something very much like the distinction he makes) to be the right one, I take his particular view to be untenable. In what follows I will first develop two independent lines of critique aimed at showing these problems. Then, in sections 4.2 and 4.3, I will sketch an alternative account that makes use of the general architecture of Sosa's epistemology. The alternative account that I defend will avoid the problems that threaten Sosa's view while offering the resources needed to answer a number of challenges that threaten internalist accounts of justification.

4.1.1 The New Evil-Demon Problem *Revisited*

Sosa's attempt to deal with the new evil-demon problem deserves closer attention. As we saw in Chapter 2, Sosa defends an account of actual-world reliabilism. According to the actual-world reliabilist:

Relative to our actual environment A, our automatic experience-belief mechanisms count as virtues that yield much truth and justification. Of course, relative to the demonic environment D such mechanisms are not virtuous and yield neither truth nor justification.¹¹⁹

The actual-world reliabilist goes on to argue that we can then assess the justificatory status of the beliefs of a victim of an evil demon relative to our environment. Since, relative to our environment, our automatic experience-belief mechanisms are reliable and

¹¹⁹ Sosa 2001: 144.

result in justified belief, a victim of a an evil demon who uses those same automatic experience-belief mechanisms will also have justified beliefs relative to our environment.

But, as I have already argued, this response seems to assume that our beliefs in the actual world are reliable. Surely, one might worry about just how helpful this type of response is once one acknowledges that the actual world could be an evil-demon world. If this were the case, then according to the actual-world reliabilist, our beliefs (in the actual world) would not be reliable and would presumably not be justified. But then the same would be true of another person who is a victim of an evil demon in a different world. Her beliefs would be unjustified relative to the actual world (our world) as well.

To his credit, Sosa attempts to deal with this kind of objection to actual-world reliabilism. Sosa could be interpreted as trying to invoke what I take to be two separate strategies for dealing with the proposed challenge. First, Sosa tries to turn the tables on the skeptic's question by considering the following conditional:

(D) If the actual world is a demon world, then our beliefs acquired through taking our experience at face value are justified.¹²⁰

According to Sosa one must treat (D) as either a material conditional or a subjunctive conditional. If we interpret (D) as a material conditional, then the conditional is true, simply because the antecedent is false.¹²¹ We are not in a demon world, so we do not need to worry about being in demon worlds. At this point the skeptic (or the foe of actual-world reliabilism) will surely throw up her hands. It seems as if Sosa is simply uninterested in the possibility that we are victims of an evil demon, even though that possibility is the kind of thing that, by his own admission, cannot be ruled out. For those

¹²⁰ Sosa 2009: 40.

¹²¹ Sosa 2009: 40.

of us who had hoped for an account of justification that offers some sense of assurance, Sosa's disinterest in the skeptical hypothesis is unsatisfactory.

Sosa will, of course, contend that his response is perfectly in line with his view. Furthermore, he would point out that if what one wants is something like intellectual assurance of the truth of one's beliefs, his account offers the necessary resources. Unfortunately, as we will see in the next section, it is far from clear whether Sosa's account of reflective knowledge stands up to scrutiny.

The most interesting way of interpreting (D) is as a subjunctive conditional. What if the actual world *is* a demon world? If so, then it seems the actual-world reliabilist is in trouble. Oversimplifying a bit, we could say that the beliefs that I have, in this world, are not justified – they are not reliably produced. As a result, appealing to my world and my belief-forming processes will not help explain why the demon-victim in another world has justified beliefs relative to my world.

Sosa's response to the subjunctive interpretation of (D) is nuanced, to say the least. When we imagine the actual world to be a demon world, we are (according to the actual-world reliabilist) taking ourselves to the closest possible world that is a demon world. But this does not mean that the actual world *is* a demon world. We are, after all, talking about a different world when we consider the demon world. As a result, we can still explain why beliefs produced in the demon world are justified relative to the actual world – the non-demon world.

The actual-world reliabilist's move here is somewhat evasive. Sosa seems to be just refusing to answer the question that the advocate of the new evil-demon problem is trying to ask. Perhaps, then the question can be asked in a different way – one that avoids

conditionals so as to stop the game Sosa is trying to play. Consider the following propositions:¹²²

1. I have had a phenomenally robust life of sensations (I have had all the same sensations that you have had).
2. My sensations have all been caused by an evil demon.
3. I form beliefs about the external world (the same beliefs you form about the external world).

and two possibly entailed propositions:

4. I have epistemically justified beliefs.
5. I have epistemically unjustified beliefs.

Now, we might ask, do propositions 1–3 entail proposition 4 or 5? When the new evil-demon problem is put this way, and the reliabilist (like Sosa) is forced to embrace either 4 or 5, the deliberate avoidance of the real problem is even harder. As we have already seen, there seems to be a powerful inclination to think that 4 is entailed by 1–3. But isn't the reliabilist committed to saying that 5 is must be true? If this line of reasoning is right, then the essence of the new evil-demon problem can be retained without invoking any messy subjunctive conditionals.

An alternative strategy Sosa might try to develop relies on his requirement of an adequate epistemic perspective. Consider a case in which a person has the coherent perspective of her own situation required for reflective knowledge.¹²³ Then consider

¹²² This strategy of avoiding conditionals, and the following propositions, come from conversations with Richard Fumerton.

¹²³ The present response to the new evil-demon problem is one that Sosa hints at, but does not seem to develop fully. In fact, on Sosa's view, it would seem to be a stretch to even

whether that perspective could remain adequate even if the person was the victim of an evil demon. If such a situation is possible, couldn't Sosa argue that the victim of an evil demon has justified beliefs because of the internal coherence (which she recognizes) of her beliefs?

While such a move might seem possible for Sosa, I am not sure it is. Remember, on Sosa's account, it is not the *awareness* of coherence alone that justifies belief and constitutes reflective knowledge. If it was, Sosa would best be characterized as an internalist. Instead, as we have already seen, Sosa argues that one builds from animal knowledge to reflective knowledge. Thus, reflective knowledge requires not only the appropriate epistemic perspective, but also requires already having animal knowledge – belief that is apt. However, by hypothesis, the person in the above case does not have apt belief. Since the beliefs in question result from processes that are not reliable, the beliefs fail to be apt.

4.1.2 The Problem of Easy Knowledge *Revisited*

Sosa's view seems attractive because it is aimed at offering an account of knowledge that can satisfy our philosophical curiosity (through reflective knowledge) while also making sense of what we commonsensically take ourselves to know (through animal knowledge). As we saw in Chapter 2, the problem of easy knowledge calls our attention to the illegitimacy of bootstrapping. There would be a problematic kind of circularity if I used my perceptual faculties to come to know that P (and that Q, and that

coherently develop the case: having the appropriate coherent perspective on one's situation seems to require "seeing" that one's automatic experience-belief mechanisms are functioning properly. However, this is exactly what the present hypothesis is trying to reject. Someone can have the appropriate coherent perspective even in the case where her automatic experience-belief mechanisms are not functioning properly.

R) and then used my knowledge that P (and that Q, and that R) to inductively infer that my perceptual faculties are reliable.

Sosa admits that this kind of blatant bootstrapping is problematic: “We cannot hope to provide a faculty with its required epistemic standing just by drawing the conclusion that it is reliable from a track-record argument based exclusively on data acquired through trusting that very faculty.”¹²⁴ In order to avoid this kind of problem, Sosa develops two theses. The first is that the epistemic standing of our animal competencies can be attained without the formation of any meta-beliefs about those competencies. The second is the transcendental argument against skepticism (or for the reliability of our cognitive faculties).

With regard to the first thesis, Sosa argues that, “our trust in our animal epistemic competencies is a source of epistemic standing for the beliefs thus acquired. This is because those competences themselves, those animal faculties, have a proper epistemic standing of their own.”¹²⁵ Thus, according to Sosa, we need not prove that the faculties are reliable in order to trust in their reliability; instead, the “default setting” is one of proper epistemic standing.

But one might still wonder how we move from the animal level to the reflective level, while avoiding vicious circularity. According to Sosa, the key is to admit that this kind of move cannot be made without circularity, and to maintain that the circularity that results is not vicious. We can exercise our animal competencies without appealing to reasons for belief (that is, the animal competencies are non-reason-involving). We can

¹²⁴ Sosa 2011: 140.

¹²⁵ Sosa 2011: 149.

use these basic competencies to support reason-involving competencies by way of an appeal to coherentism:

Reflective endorsement may now take its place in the web with no apparent special problems. Through our growing knowledge of ourselves and the world around us and of the relation between the two, we come to see our modes of rational basing and other belief acquisition as sufficiently reliable.¹²⁶

As we saw in Chapter 2, Sosa's appeal to coherence invites its own problems.

Certainly it is not the case that the mere fact that one's beliefs cohere results in the kind of reflective justification Sosa is attempting to characterize. Sosa makes clear that what is needed is some kind of endorsement of the coherence in question. Again, something like a disposition to endorse is not enough either, for if it were, then animal knowledge and reflective knowledge would not be importantly different from one another. Thus Sosa owes us an account of reflective knowledge that makes sense of something like an access condition on the coherence that is required for reflective knowledge.

Furthermore, as I have already argued, if Sosa retreats from his insistence on the condition of coherence as being a necessary component of reflective knowledge, it is not clear why reflective knowledge is any better than animal knowledge. In fact, on this account, reflective knowledge just is *more* animal knowledge. As I showed in Chapter 2, although Sosa does (at times) seem to think of reflective knowledge as being just more animal knowledge, this weak characterization of reflective knowledge fails to live up to Sosa's promise of a philosophically satisfying account of knowledge.

Sosa insists that in the end, no non-circular account of justification can be given. As a result, he argues, the circularity entailed by his account seems no worse than that of

¹²⁶ Sosa 2011: 151.

the paradigm internalist hero, Descartes. However, as I will show in section 3 of this chapter, this supposition is false.

4.2 The Apparatus of Epistemic Descent

While I take the problems developed in the preceding section to ultimately leave Sosa's epistemology implausible, I think there is something essentially right about his distinction between different *kinds* of knowledge (and justification). In what follows I will sketch what I take to be an account of justification that clearly borrows from Sosa's work. However, because of a crucial methodological difference in how the distinction between kinds of justification is developed, the view does not fall victim to the same kind of problems that plague Sosa's account.

In Chapter 2 I argued that Sosa's account of knowledge is best understood as being a bottom-up account: one starts with an epistemologically undemanding concept of knowledge, and then, by adding further epistemic elements, works one's way up to a more philosophically satisfying kind of knowledge. As Sosa might say, one achieves epistemic ascent by "climbing the epistemic ladder." In contrast, the view I develop is best understood as a top-down account. One starts with a paradigm conception of justification – ideal justification – and then develops non-ideal kinds of justification by stripping away the layers (or components) that constitute ideal justification. In this way, as someone strips away the components required of ideal justification, one is moving down the epistemic ladder. The account I offer is thus an account of epistemic descent. If an account of justification does not satisfy the requirements on ideal justification, then by definition, it is less than ideal, which is what I will call *degenerate* justification. However, as I will show, the accounts of degenerate justification that I aim to explore

have something important in common with ideal justification, and this is why they are still properly understood as being kinds of justification, rather than degrees of justification.

Some clarification of the apparatus of epistemic descent is in order. One could develop an account of different degrees of justification according to which ideal justification and degenerate justification are at opposite ends of a spectrum. On this view, ideal justification might be considered the upper limit of justified belief – perhaps justification that entails certainty. Degenerate justification, on the other hand, could be construed as something less than ideal or certain, but that results in highly probable belief. This kind of distinction can be illustrated with a lottery: if I held all the tickets in a lottery, we would say my chances of winning were certain, and thus my justification for the belief that I held the winning ticket was ideally justified. Alternatively, if I held the majority of the tickets, but not all of them, we might say I was degenerately justified in my belief that I held the winning ticket. Yet since I did not have all the tickets, my belief would lack certainty and would not reach the level of ideal justification.

This scenario does not describe my project. The distinction that I am developing is not aimed at distinguishing between *degrees* of justification, but (like Sosa) between *kinds* of justification.¹²⁷ Thus, if one of the constituent parts required for ideal justification is not present, then by definition, one would not have ideal justification. However, if the account I develop is successful, then the remaining criteria (while insufficient for ideal justification) might be sufficient for degenerate justification. On my view, ideal justification and degenerate justification are intimately related, but the latter

¹²⁷ Of course, as I have noted before, Sosa's primary focus is on kinds of knowledge, whereas mine will be on kinds of justification.

lacks some significant epistemic element. Consider, as an example, the necessary and sufficient conditions for certain geometrical figures. An equilateral triangle can be defined as a three-sided figure with three equal angles and three equal sides. If the sides of the triangle are not all equal, or the angles are not all equal, the figure fails to be an equilateral triangle. The triangle with unequal sides or unequal angles is still a triangle, but it no longer meets the necessary and sufficient conditions for being an equilateral triangle. On my view, ideal justification is like the equilateral triangle, while degenerate justification is like the scalene or isosceles triangle. They are all *kinds* of triangles (justification), just not the same kind.

4.2.1 Epistemic Descent and Noninferential Justification

In my presentation of noninferential justification¹²⁸ I intentionally ignored a distinction that now must be brought to the forefront. It is helpful in discussions of justification to distinguish between what has come to be called *propositional* justification and *doxastic* justification. To see the importance of this distinction, imagine the following case:

Leah sees a storm on the horizon, she hears a severe-weather warning being broadcast through her weather radio, and she remembers from the morning news that there is a 100% chance of rain today. Leah forms the belief that it will rain today.

It seems like we would want to say that Leah's belief that it will rain today is justified. Leah has what appears to be very good evidence: she sees the storm coming, she is aware of a severe-weather alert for her area, and she remembers the weatherman confidently predicting rain for that day. However, even if Leah has very good evidence for the belief that it will rain today, it seems like her belief must also be *based on* that

¹²⁸ See section 3.3.

evidence for the belief to be justified. If Leah believed that it will rain today because she had just asked a magic 8-ball, “Will it rain today?” and the magic 8-ball’s response was, “All signs point to yes,” we might be more hesitant to say that Leah’s belief is in fact justified.

The distinction between propositional justification and doxastic justification allows us to adequately deal with cases like Leah’s 8-ball reasoning. Leah’s belief that it will rain today is propositionally justified because she *has* sufficiently good evidence to justify her belief. Crucially, though, the evidence that allows us to say that Leah’s belief is propositionally justified is that she sees the storm coming, she is aware of a severe-weather alert for her area, and she remembers the weatherman confidently predicting rain for that day. However, as I developed the case, this evidence is not what Leah *used* as the basis of her belief. Instead, Leah based her belief on the answer from the magic 8-ball. Since Leah did not base her belief on the evidence that propositionally justified her belief, the belief was not doxastically justified.

In order for Leah’s belief to be doxastically justified, the basis of Leah’s belief must be the evidence that makes the belief propositionally justified. Because an account of doxastic justification depends on an account of propositional justification, it is natural to start with an analysis of propositional justification and then add the further condition required for doxastic justification. As the previous example was intended to illustrate, this further condition is an appropriate basing relation that holds between one’s evidence and

one's belief. Furthermore, it seems to me that this basing relation is best understood as a causal relation.¹²⁹

In addition to thinking that propositional justification is a more fundamental concept than doxastic justification, I think that propositional justification is of more interest to the philosopher, *qua* philosopher, than doxastic justification. As we saw in Chapter 3, the most philosophically interesting way of carving up the internalism/externalism debate had to do with the sense of assurance that accompanied internalist conceptions of justified belief. The internalist's project seemed to be one of doing epistemology from the first-person perspective.¹³⁰ As I reflect, from the first person perspective, on which of my beliefs are justified, the question is concerned with propositionally justification: I am trying to figure out which of my beliefs have evidence that would justify those beliefs. This is a question that can be answered from the armchair. However, I may simply not be in a position to know the causal origins of my beliefs from the first-person perspective.¹³¹

¹²⁹ This understanding is by no means uncontroversial. Huemer seems to think it involves a special sorting of 'seeing,' Fales thinks it is a *sui generis* connection that 'moves' one from acceptance of premises to acceptance of a conclusion, but where the 'movement' is not causal.

¹³⁰ The internalist's preoccupation with first-person epistemology, as I argued in section 3.2, can be contrasted with the externalist's preoccupation with third-person epistemology.

¹³¹ Of course, the plausibility of this claim will depend a lot on the account of causation one accepts. If one is a generality theorist and one thinks the relevant generalizations (that are partly constituent of) are known empirically, it is not clear how anyone would know what the causes of belief are. People brought up as Protestants tend to be Protestants – this is no accident. Many people might try to deny that it was their parents that caused them to have the beliefs they do, but it is hard to avoid the conclusion that they were causally influenced. Certainly the parents' action does enter into a causal explanation of what they believe.

Because of these considerations, my conceptions of ideal justification¹³² are intended to be understood as conceptions of propositional justification. However, each of the accounts of degenerate justification that I develop will involve a basing relation, and as a result, they are best understood as characterizing doxastic justification.

To see how a theory of degenerate noninferential justification can be developed, we need to consider, again, the account of ideal noninferential justification presented in the previous chapter:

INJ: S has ideal noninferential justification for the belief that P when:

1. S has the thought that P,
2. S is acquainted with the fact that P, and
3. S is acquainted with the relation of correspondence holding between the thought that P and the fact that P.

As condition 3 makes explicit, being able to *see* the connection (or the relation of fit) between the thought that P and the truth maker for that thought (the fact that p) allows one to be ideally justified in the belief that P. However, we can conceive of an account of degenerate noninferential justification that fails to meet the criteria set by INJ:

DNJ: S has degenerate noninferential justification for the belief that P when:

1. S has the thought that P,
2. S is acquainted with the fact that P, and

¹³² See section 3.3 for my account of ideal noninferential justification and 4.2.2 for my account of ideal inferential justification.

3. S's belief that P is caused by the fact that P (where the fact that P is the truth maker for the belief that P) and there is relation of correspondence that holds between the fact that P and the belief that P.

Stripping away from ideal noninferential justification the condition that S must see the connection between the belief and its truth maker (i.e., the connection of correspondence that holds between the thought that P and the fact that P) results in a kind of justification that is clearly short of ideal. Because there is a relation of correspondence between the fact that P and the belief that P, and because the belief that P was caused by the fact that P, we can still say that S's belief that P is justified, albeit degenerately.

Why, then, is condition 3 from INJ required for justification if one can still have justified belief (as in DNJ) while lacking that condition? The answer is simple: in the case of INJ, one has before one's consciousness everything needed to see that the belief is justified, and this is why (as I argued in Chapter 3) the account of ideal noninferential justification results in a sense of assurance: one can *see* what one's belief has going for it. The account of degenerate justification offered in DNJ does not require that one see the connection between the belief and its truth maker. However – and this is crucial – the belief is only degenerately justified if there is a causal relation between the belief and the truth maker that (in some sense) mirrors the non-causal relation required in condition 3 of INJ. Even though DNJ does not require that one see the connection between the belief and the truth maker, it does require that there be an appropriate connection between the belief and the truth maker.

Since it seems that experience can contain a very large number of complex facts, one might worry that my account of DNJ justifies belief too easily. Consider, for example, the case of the speckled hen:

Let us consider the visual sense-datum which is yielded by a single glance at a speckled hen. ...Our problem pertains to the question: how many speckles does the datum comprise?¹³³

Let us grant that I am having an experience of a speckled hen, and that the speckled hen has forty-eight speckles. If I find myself believing that the hen has forty-eight speckles, wouldn't DNJ entail that my belief is justified?

Since DNJ requires that one be acquainted with the fact that P in order for the belief that P to be justified, I think the appropriate answer to this question hinges on whether or not we are acquainted with every fact exemplified in experience. Since I do not think we are acquainted with every such fact, I am inclined to suggest that in the case of the speckled hen we are simply not acquainted with the fact that the hen has forty-eight speckles, and this is why my belief about the speckled hen would fail to be degenerately (or ideally) justified.¹³⁴

The lesson to learn here is that some fact's being exemplified in one's experience does not imply that one is acquainted with that fact. Since DNJ requires that one be acquainted with the truth maker (the fact) for one's belief (the truth bearer), it is clear that justification does not come too easily (or not as easily as it may have seemed).

¹³³ Chisholm 1942: 368.

¹³⁴ Alternatively, one might argue that DNJ requires some further condition of having an ability to see the relation of correspondence given enough time, practice, training, etc. As such, we could not become aware of the relation of correspondence, and as a result, the case would not fit the additional condition on DNJ. I am less enthusiastic about requiring this further condition.

4.2.2 Fumerton's Principle of Inferential Justification

As I argued in Chapter 3, the account of noninferential justification as I have developed it is only the base clause of a recursive analysis of justification. In order to account for inferential justification and to complete the recursive analysis of justification, the development of a recursive clause is in order. As was the case with the account of noninferential justification, the account of inferential justification initially offered will constitute what I take to be ideal inferential justification.

Suppose that my pregnant wife tells me that we need to start purchasing new clothes for the baby that is due in eight months. When I ask what is wrong with all the perfectly good baby clothes we already have, she responds by saying that those are all girl clothes, and she believes we are having a boy. Should I believe that (P) I need to purchase new clothes for the baby, because my wife believes that (E) we are having a boy? One might wonder why my wife is convinced that she is having a boy this early in her pregnancy. If I find out that my wife's belief that she is having a boy is based on a mere hunch, do I have an epistemic reason to believe that I should purchase new clothes for the baby? No.

If one is to get justification from some other belief, that belief must itself be justified. From these types of cases, we can conclude that in order to be justified in believing P on the basis of E, one must at least be justified in believing E. But is this enough? Suppose I believe that I am going to have a long and rewarding life. When friends ask why it is that I believe I am going to have a long and rewarding life I tell them it is because I consulted a magic 8-ball. It may very well be true that I have excellent justification for my belief that a magic 8-ball indicated that I am going to have a long and

rewarding life, but something still seems to be missing. Friends might further ask why it is that I think my belief that an 8-ball said I will have a long and rewarding life makes it all likely that I will. In this case, the problem with the inferential justification is not that my belief that P is based on an unjustified belief that E, but that I have no reason for thinking that E would make it probable that P.

On my view, ideal inferential justification requires satisfaction of what Fumerton has called the principle of inferential justification:

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.¹³⁵

It is important to note the skeptical implications that this analysis of inferential justification seems to invite. Attempts to account for the kind of justified belief required in clause 2 have a long and varied history in modern philosophy.¹³⁶

The kind of argument that I have given in favor of clause 2 of Fumerton's principle of inferential justification has come under attack by Michael Huemer.¹³⁷ To see the problem that Huemer is proposing, consider his case of an astrologer:

Suppose one ran into an astrologer who claimed to be justified in believing that there will be many wars in the year 2000 and who offered as his evidence that Jupiter will be aligned with Saturn that year. We might challenge his justification for believing that Jupiter will be aligned with Saturn in the year 2000, but we would be much more likely to focus on the connection between the position of the planets relative to one another and the behavior of people here on earth. In short, we would challenge his claim to have justified belief in the one proposition (P) on the basis of the

¹³⁵ Fumerton 1995: 36.

¹³⁶ Descartes's use of God as a non-deceiver, Hume's attack on inductive inference, and efforts to rationally reconstruct inference to the best explanation come to mind.

¹³⁷ Huemer 2002.

other (E) by questioning his justification for believing that E makes probable, confirms, makes likely to be true P.¹³⁸

According to Huemer, the astrologer who infers predictions about wars in 2000 from facts about the location of heavenly bodies does not make such predictions on the basis of the location of the heavenly bodies *alone*. This type of reasoning (including my 8-ball example) is enthymematic:

The astrologer does not merely believe [that Jupiter will be aligned with Saturn in the year 2000]; she has a host of background beliefs about the characteristics of the planets, how they influence human affairs, and the like. As a simple example, the astrologer might believe [that the alignment of Jupiter with Saturn causes people to be more hostile,] and then it would be the conjunction of [her beliefs] that the astrologer would base her belief [that there will be many wars in the year 2000] on.¹³⁹

Huemer's claim is that once the further premises or beliefs are made explicit, the reasoning does not seem so problematic (although, the explicit background premise will still strike us as problematic and this will explain our reservations about the argument). The advocate of the principle of inferential justification could admit that the examples in question might be best construed as being enthymematic while arguing that the essential idea being exemplified still stands.

In the case in which one's evidence entails a conclusion, unless one is able to see the connection between the evidence and the conclusion, the advocate of the principle of inferential justification contends that something still seems to be missing. Consider a case involving a simple set of mathematic truths. Using an appropriate such set, the skilled

¹³⁸ Fumerton 1985. Fumerton 2006 also discusses an example of a palm reader. Imagine a palm reader who bases the belief that you will have a long life on her justified belief that you have a long "life line" on the palm of your hand.

¹³⁹ Huemer 2002. I have inserted text to help keep Fumerton's example and Huemer's consideration of the example consistent.

mathematician could prove an impressive array of complex mathematic theorems. It is plausible to think that the skilled mathematician can see the truth of the simple mathematical truths as well as that of the complex truths. In other words, the skilled mathematician can see the connection between the premises and the conclusion.

If we imagine my own comprehension of a set of simple mathematical truths, it also seems plausible to say that I can see that they are true. However, were my belief in the complex mathematical truths to be questioned, it would be clear that I do not have anything like the justification that the skilled mathematician has. Both of us have evidence that entails the same conclusion, but only one of us is able to see the relation between that evidence and the conclusion. As a result, only one of us has inferentially justified belief, according to the principle of inferential justification. My justified belief in some set of simple mathematical truths, along with the fact that those simple mathematic truths entail some other complex mathematical truth, do not alone result in a justified belief about the complex mathematical truth. On the view I am advocating, unless I am able see the connection, I am not inferentially justified.

The conceptions of degenerate inferential justification that I will be developing are parasitic upon the account of ideal inferential justification that I defend. As we will see, Fumerton's principle of inferential justification is the basis for my account of ideal inferential justification. If it could be shown that ideal inferential justification is impossible, my accounts of degenerate justification would be undermined.

As we have already seen, clause 2 of Fumerton's principle of inferential justification requires that one be justified in believing that E makes probable P in order to be justified in believing that P on the basis of E. However, one might wonder how

justified belief in this kind of evidential connection (i.e., my belief that E makes probable P) is achieved. On my view, evidential connections that hold between E and P do so necessarily and are instances of *a priori* truths.¹⁴⁰ In effect, the view I am advancing is that evidential connections can be *seen*. It is the seeing of these evidential truths that justifies the beliefs in question when the relevant evidential relations obtain.

Opponents of the principle of inferential justification and my account of ideal inferential justification might argue that such an account leads to a vicious regress. Consider Lewis Carroll's dialogue between the Tortoise and Achilles,¹⁴¹ where Achilles is trying to prove that

(Z) The two sides of this triangle are equal.

Achilles argues that (Z) follows logically from two premises:

(A) Things that are equal to the same are equal to each other, and

(B) The two sides of this triangle are things that are equal to the same.

However, the Tortoise is able to get Achilles to accept

(C) If (A) and (B) are true, then (Z) is true.

But once Achilles accepts that (C) is required in order to prove (Z), he is committed to an infinite regress. The Tortoise claimed that (A) and (B) by themselves do not imply (Z); one needs the further premise (C) that links (A), (B), and (Z). However, if (A) and (B) weren't sufficient to infer (Z), then (A), (B), and (C) will not get the job done either. In addition, says the Tortoise, Achilles needs to accept

(D) If (A), (B), and (C) are true, then (Z) is true.

¹⁴⁰ That these evidential connections hold necessarily is of crucial importance to my view. See section 5.1.3 for an explanation of why this is the case.

¹⁴¹ Carroll 1895.

Of course, (D) does not stop the regress. In addition to (A), (B), (C), and (D), Achilles is committed to

(E) If (A), (B), (C), and (D) are true, (Z) is true.

The lesson for Achilles is that once he grants that (A) and (B) are not sufficient to imply (Z), he will never be able to finish writing down the required premises in order to prove (Z). What Achilles should have realized is that one does not need some further premise to strengthen an argument in which the premises entail the conclusion. While it is reasonable to say that Achilles must be able to see the connection between the premises and the conclusion in order to justifiably believe the conclusion, this role that seeing plays had better not amount to involving some further premise – one in an infinite series of premises required in order to make the necessary inference. In order to avoid the kind of vicious regress that Carroll's dialogue invites, the advocate of ideal inferential justification must be able to understand the connections between premises and conclusion in a way that does not ultimately depend on an infinite series of further inferences.¹⁴²

Of course, my goal is not to explore the limits of ideal justification (i.e., what beliefs can be justified when each condition on ideal justification is met). Instead, my goal is to explore the possibility of degenerate justification (i.e., a different kind of justification that fails to meet the standards on ideally justified belief).

4.2.3 Epistemic Descent and Inferential Justification

How might the apparatus of epistemic descent be employed in the context of a discussion of inferential justification to capture some sense of rational belief? The

¹⁴² I will explore how this challenge might be met in section 5.2.

account of degenerate justification I will develop depends on a relation between certain psychological states (justified occurrent beliefs, sensations, memories, and justified dispositional beliefs) and the beliefs those psychological states *cause*. We can contrast the account of ideal inferential justification with different (and increasingly degenerate) conceptions of degenerate inferential justification.

Ideal Inferential Justification (IIJ): A subject S has ideal inferential justification for the belief that P on the basis of E when:

1. S believes that P,
2. S justifiably believes that E, and
3. S justifiably believes that E makes probable P.

Examples of the kind of ideal inferential justification captured by IIJ can be seen in instances of justified belief that involve clear inferences. For example, if I have the justified noninferential belief that P, I could infer that either P or Q. Here my ideally justified inferential belief is the result of a noninferential belief, my recognition of truths about disjunctive propositions, and my awareness of the entailment.

The type of reasoning employed in this example may seem a bit odd, especially if it was supposed to illustrate a paradigm case of inferentially justified belief. A more complicated example, and one that is perhaps more representative of many of our ideally justified inferential beliefs, is the Monty Hall problem:

Imagine that you are on a game show and you are given the chance to win a prize. There are three closed doors, behind each of which there is a randomly-placed prize out of your view, and you are told you will receive the prize behind whichever door you select. Two of the doors have a goat behind them (which you don't want) and one has a car behind it (which you do want). You pick a door, but before that door is opened, the host of the game show (who can see behind the doors) opens one of the doors you did not pick and reveals a goat. After the door is opened, the host says you can change your selection.

Should you stay with your first choice, or should you switch your choice to the other remaining door?¹⁴³

Initially it seems as if switch would not be beneficial: your original choice has a 1 in 2 chance of hiding the car, just like the other remaining door. But this thought is wrong. Your initial choice had a 1 in 3 chance of being correct. Nothing that the game show host does affects this initial probability. Thus, your chances of winning if you don't switch your selection are 1 in 3. However, if you switch, your chance of winning the car are 2 in 3.

Why do I think this example is such a helpful example of ideally justified inferential belief? My belief that the best option is to switch to the other door is clearly inferential. It is something that I had to be shown to be true – initially the switch did not seem beneficial, but when the solution of the case is explained, I can see why the right answer must be right. The inferred belief depends on my other beliefs about the case, and it depends on my beliefs about probability. The fact that I am able to see how the truths of probability come to bear on this particular case is crucial to its being a case of ideally justified inferential belief.

In the simplest case of degenerate justification, S is not aware of the epistemic relation that holds between E and P (that is, S fails to see the connection that is required for justified belief in condition 3).

Degenerate Inferential Justification (DIJ): A subject S has degenerate inferential justification for the belief that P on the basis of E when:

1. S believes that P,

¹⁴³ This revised case is based on a problem presented by Steve Selvin in a letter to the editor of *American Statistician*.

2. S justifiably believes that E, and
3. S's belief that P is caused by E* (where E* is the truth maker for E) and there is an evidential relation holding between E and P (E makes probable P).

On this account of degenerate justification, it is S's experience of E* – the truth maker for E – that causes S's belief that P. Since E* is the cause of P, and E makes probable P, S's belief that P is degenerately justified. Of course, one might wonder how often we satisfy the conditions of DIJ. An answer to this question will (at least) involve metaepistemological questions about justification.

For example, according to phenomenal conservatism, S has some degree of justification for the belief that P if and only if it seems to S that P.¹⁴⁴ If DIJ were true and we were in an experiential state such that it seemed to us that P, that is all that we would need for our justified belief that P.¹⁴⁵ While Huemer might object to the labeling of his view as an instance of degenerate justification, he would certainly not object to the idea that one can have a justified inferential belief without seeing the evidential connection that holds between the truth maker and the belief. According to a phenomenal conservative like Huemer, the vast majority of our justified beliefs could be characterized by the kind of justification outlined in DIJ.

However, if we are concerned that phenomenal conservatism won't stand up to scrutiny – that such a view fails to adequately stand up to skepticism – we might have to

¹⁴⁴ For the clearest and most contemporary versions of this kind of view, see Huemer 2001 and 2007.

¹⁴⁵ I am setting aside plenty of further considerations here. For example, Huemer argues for a no-defeaters clause, and has more recently argued that a phenomenal seeming is not enough for *prima facie* justification. Instead, a phenomenal seeming results in some degree of justification.

develop further (and more degenerate) conceptions of degenerate justification. If, for example, one were more sympathetic to the idea that the majority of our beliefs result from some complex inference to the best explanation, DIJ (which only offers an account of degenerate justification where S's belief is caused by the truth maker for E) will still be too restrictive.

A more complicated case is one in which the truth maker for E is not just between E and P, but one that holds between a much more complex set of propositions E^1 – E^n and P:

DIJ*: A subject S has degenerate inferential justification for the belief that P on the basis of E when:

1. S believes that P,
2. S is in psychological state E^* (where psychological state E^* includes a complex set justified occurrent beliefs, sensations, apparent memories, and justified dispositional beliefs), and
3. S's belief that P is caused by E^* and there is an evidential relation holding between a complex set of propositions that describe psychological state E^* and P (E^* makes probable P).

On this conception of degenerate justification, what is causing S to believe that P is not simply the *lone* truth maker for P. Instead, the cause of the belief that P is some *set* of S's psychological states E^* (S's sensations, justified dispositional beliefs, etc.) which together constitute the truth makers for propositions that, when taken together, make P probable. According to DIJ*, degenerate inferential justification requires that evidential

connections obtain between propositions *made true by* S's psychological state and the belief in question, but that S does *not* necessarily believe.¹⁴⁶

Finally, once we realize that the causal connection exhibited in DIJ* is between a complex psychological state that includes the agent's vast array of background beliefs, and we can further imagine an even more degenerate conception of justification according to which an agent has as part of the psychological states some part of the evidence that would justify the belief that P, but does not have all of such evidence.

We can capture the difference between DIJ* and an even more degenerate conception of justification (DIJ**) by considering the difference between the causal connections that would exist in a child and those that would exist in an adult. We can imagine that an adult, through her years of experience and her vast array of background beliefs, has a psychological state complex enough, with enough evidence, that it makes probable the inferred belief. A young child's psychological states have some degree of complexity, but they lack the complexity of the adult's. As a result, while the child's psychological state might contain some of the evidence that could be used in an inference, it is reasonable to suppose that she does not possess all of the required evidence to make a legitimate inference, and as a result, she fails to meet the criteria set forth in DIJ*.

DIJ**: A subject S has degenerate inferential justification for the belief that P on the basis of E when:

1. S believes that P,

¹⁴⁶ Fumerton 2004 makes a similar point when he characterizes unreflective justification.

2. S is in psychological state E^{**} (where E^{**} includes a set of justified occurrent beliefs, sensations, apparent memories, and justified dispositional beliefs), and
3. S's belief that P is caused by E^{**} and E^{**} is part of the evidence that would be required for the evidential relation that would hold between a complex set of propositions that describe a psychological state E^* and P (E^* make probable P).

4.2.4 Clarifications

Justified dispositional belief plays a crucial role in my account of degenerate justification. I take dispositional beliefs to be beliefs that one has – propositions that would be occurrently believed if they were entertained. Dispositional belief, then, must be distinguished from a disposition *to believe*. One might be tempted to defend an analysis of dispositional belief that amounts to the following subjunctive conditional: *if S were to consider some proposition P, then S would believe P*. However, this analysis is vulnerable to some obvious counterexamples. For example, if Richard had been knocked unconscious, then it is still true that (on the subjunctive analysis of dispositional belief) he would believe that I was awake *if* he were to wake up and be asked whether I was awake. Yet attributing dispositional beliefs to an unconscious person seems a stretch.

I am inclined to think that these kinds of counterexamples to the subjunctive analysis of dispositional belief should persuade us that there is more to dispositional belief than merely what can be captured by the subjunctive conditional. One might think, for example, that in order to distinguish dispositional belief from a disposition to believe, one could require of dispositional belief that the belief in question must have once been

believed occurrently, and perhaps that there must be an occurrent ground that can be traced to that dispositional belief. However, this more subtle analysis runs into its own problems. It seems that we would want to say that minutes ago I dispositionally believed that $1+1=2$. We would also probably want to say that I dispositionally believed that $2+1=3$, and so on – eventually we will get to some simple arithmetical belief that I have never occurrently believed. Is the present account correct to preclude this belief from being one of my dispositional beliefs?

We could try to fix the problem by dropping the condition on dispositional belief that requires one to have believed each dispositional belief occurrently. Instead, it might be argued that an account of dispositional belief depends on what is justified given our present experience and *would* be immediately believed if considered. Yet this seems to return us dangerously close to the subjunctive analysis.

Despite my lack of a refined notion of dispositional belief, it is clear that if a dispositional belief is to be justified, it needs grounding. An analysis of justified dispositional belief, like occurrent belief, will involve an analysis of both inferentially justified dispositional belief and noninferentially justified dispositional belief.

In the case of noninferentially justified dispositional belief, an adequate grounding might be an experiential state that causes one to dispositionally believe some proposition when there is an evidential relation that holds between propositions that describe the experiential state and the dispositional belief. For example, before I was consciously attending to it, I had the justified dispositional belief that I seem to hear a fan blowing behind me. The noninferential dispositional belief that results from this experiential state is grounded because one of the propositions that describes the

experiential state is that *I seem to hear a fan blowing behind me*, and crucially, there is an evidential relation that holds between the proposition describing the experiential state and the dispositional belief. In the case of inferentially justified dispositional belief, an adequate grounding might include other justified dispositional beliefs. For example, a few moments ago I had the inferentially justified dispositional belief that tired children are cranky. This belief was justified in virtue of its evidential relation to other dispositional beliefs, such as that most tired children I have dealt with have been cranky.

Ultimately, our dispositional beliefs, apparent memories, and the like are justified in virtue of the fact that there are evidential relations that hold between the propositions describing the causes of belief and the propositions which describe our vast arrays of background beliefs. And, to be clear, awareness of these evidential relations is simply not required on the conception of degenerate justification I am offering.

Before I move on to a discussion of how the apparatus of epistemic descent can be utilized in an explication of some of the problems that threaten internalism, it may be necessary to further articulate why it is that I take the accounts of degenerate justification to be understood as not only a kind of justification, but as a kind of internalist justification. After all, the conceptions of degenerate justification that I have developed have some obvious externalist elements. Specifically, I appeal to a causal relation that, by definition, the subject is unaware of, yet I still contend that the belief in question is justified (albeit degenerately) in an internalist sense. How do I reconcile my reliance on a causal relation with the claim that degenerate justification is *still* internalist justification?

Essentially, the accounts of degenerate justification that I have developed still require the same evidential relations that are required by ideal justification. Ideal

justification requires that the appropriate evidential relations exist and that one is aware of those evidential relations. Degenerate justification still requires that the appropriate evidential relations exist, but substitutes an appropriate causal relation in place of the awareness required by ideal justification. My claim is that the epistemic work is done by that very same evidential relation.

CHAPTER 5

THE DIVIDENDS OF EPISTEMIC DESCENT

5.1 Results from the Apparatus of Epistemic Descent

In what follows I will begin the process of showing how one can use the resources of epistemic descent to address many of the problems that have threatened internalism. The account of justification I develop (which includes both ideal and degenerate kinds of justification) is aimed at making sense of how many of the beliefs we commonly take to be justified are in fact justified while also taking skepticism seriously.

5.1.1 Goldman's "Internalism Exposed"

As we saw in Chapter 1, Goldman was concerned that the internalist places too high a requirement on justification. The implication of this lofty requirement, Goldman argued, was skepticism with regard to many of the things we commonsensically take ourselves to know. My defense of an account of ideal inferential justification has made apparent how complex the justification of commonsensical belief really is. Given this complexity, is Goldman correct in saying that the internalist is hopelessly wedded to skepticism?

In what follows I will consider each of Goldman's arguments against internalism presented in the first chapter. I will use the account of degenerate justification developed in section 2.1 to show why, in some cases, the internalist need not embrace skepticism. However, I will argue that while the apparatus of epistemic descent seems to offer very promising resources to help the internalist avoid skepticism in some cases, in other cases a more radical approach might be necessary in order to make sense of justified belief. To be clear, the essential move I will make in each of the following responses will be to

highlight the role that one's psychological states can play in degenerately justified belief.¹⁴⁷

5.1.1.1 Problems of Unavailable Evidence *Revisited*

Goldman first considers two problems that deal with *unavailable* evidence. In the first problem, the evidence in question is what we would colloquially think of as being “out-of-mind.” As Goldman puts it,

At any given time, the vast majority of one's beliefs are stored in memory rather than occurrent or active. ...Furthermore, for almost any of these beliefs, one's conscious state at the time includes nothing that justifies it.¹⁴⁸

Goldman's claim seems accurate – most of our beliefs, including those that are used to justify another belief, are not occurrent. But how, then, can internalists make sense of their justification? The account I am defending allows for the possibility of justified dispositional belief. We have seen that ideally justified inferential and noninferential beliefs seem to require awareness of the propositions believed (and consequently that those belief states be occurrent), but the account of degenerate justification I have sketched makes no such requirement.

Many of these justified dispositional beliefs are justified in virtue of their relation to other justified dispositional beliefs.¹⁴⁹ And, as the account of degenerate justification makes clear, this relation is causal. As a result, there is no reason to think of justified

¹⁴⁷ And, as in section 4.2.1, I understand one's psychological states to include justified occurrent beliefs, sensations, and justified dispositional beliefs.

¹⁴⁸ Goldman 1999: 278.

¹⁴⁹ These other justified dispositional beliefs could themselves be justified inferentially or noninferentially.

dispositional beliefs as requiring the kind of attentiveness that Goldman seems to think characterizes internalism.

Goldman also considers a more troubling problem. While the first problem called attention to non-occurrent evidence, the second problem focuses on *forgotten* evidence: evidence that is literally out-of-mind. Goldman poses the case of Sally, who seems to remember that eating broccoli is beneficial, but can no longer recall her evidence for that belief. Furthermore, imagine that in one case Sally's evidence was reading about broccoli in the *New York Times*, and in another case, Sally's evidence was reading about broccoli in the *National Enquirer*. According to Goldman, Sally's belief in the former case is justified, while her belief in the latter case is not. He argues that the internalist cannot account for this difference.

In order to deal with this case, a few clarifications are in order. If this problem is to be truly different from the first, then we would need to say that the relevant evidence – the beliefs on which Sally's broccoli-beliefs are based – are really forgotten, which is to say she no longer has these beliefs occurrently or dispositionally. If this is the case, and all Sally has is an inclination to accept a proposition (that she seems to remember believing that eating broccoli is beneficial), then it appears Goldman is right in saying that the internalist cannot account for the difference between the cases. However, I think Goldman is wrong to say that in the one case Sally is justified, and in the other case she is not justified. The internalist should, I think, contend that in both cases Sally's belief is

unjustified (in the ideal sense). Sally, as Goldman points out, no longer has before her mind the necessary justification contributors to justify her belief.¹⁵⁰

Perhaps, though, we can construe the case as one in which Sally's psychological state includes a complex set of justified dispositional memory beliefs and the justified dispositional belief that propositional memory is generally reliable. In this scenario, if Sally's robust psychological state stands in the right kind of causal connection to Sally's belief about the health benefits of broccoli, then the account of degenerate justification might be able to offer an analysis of how Sally's belief is (degenerately) justified in both cases.

One could perhaps appeal to an entirely causal analysis of justified belief – one that involves not only the current psychological state of an agent but also the past psychological states of the agent (as Goldman does), but such a move seems so far removed from the starting point of ideal justification that I am inclined to suggest that it just does not have enough in common with the account of ideal justification to still warrant the locution of “justification.”

5.1.1.2 Problems of Cognitive Limitations *Revisited*

Goldman also considers two problems that deal with cognitive limitations. He contends that internalism is threatened by what can be concurrently retrieved and by the amount of time (the doxastic interval) required for the necessary retrieval and computation. Again, it is obvious that what Goldman has in mind is a variety of internalism that requires some kind of access to justification contributors and some kind

¹⁵⁰ Of course, my lack of enthusiasm for an adequate internalist solution to the problem of forgotten evidence depends on what I take to be the worst-case scenario. If one allows for a more robust sense of “remember,” such as *episodic* memory as opposed to *propositional* memory, I think the chances of success are much greater.

of explicit inference being made. However, as we have seen, the account of degenerate justification I develop does not require either of these.

Instead, the advocate of epistemic descent can argue that the account of degenerate justification collapses the issue of a doxastic interval since degenerate justification does not require the subject to recognize that the relevant justification contributors obtain. Instead, in order for someone to be degenerately justified, the crucial requirement is an evidential relation holding between the belief and the psychological states that caused that belief. Furthermore, because this process does not require the retrieval of justification contributors, the problem of concurrent retrieval does not apply.

5.1.2 The Gettier Problem

I have already shown that my account of epistemic descent has the resources to deflect many of the skeptical worries that have traditionally threatened internalist accounts of justification. In what follows, I will argue that my account of epistemic descent can also go a long way toward answering one of contemporary epistemology's nagging problems – the Gettier problem. In *The Problems of Philosophy*,¹⁵¹ Russell defends a *justified true belief* (hereafter JTB) analysis of knowledge from Gettier-style counterexamples. I will argue that Russell's solution¹⁵² to the Gettier problem, when combined with my account of epistemic descent, offers the resources necessary to dissolve many of the purported Gettier-style counterexamples.

¹⁵¹ Russell 2004.

¹⁵² Russell's solution, strictly speaking, is a solution to a problem that Russell himself anticipated in *The Problems of Philosophy* long before Gettier. However, this species of counterexample has come to be known as the Gettier problem, and I will follow suit.

Russell contends that true belief is not sufficient for knowledge, and he argues that two further conditions must be met on what can be called knowledge: “It is clear that a true belief is not knowledge when it is deduced from a false belief. In like manner, a true belief cannot be called knowledge when it is deduced by a fallacious process of reasoning...”¹⁵³ Consider Russell’s case of Mr. Balfour:

(1) I am justified in my belief that Mr. Balfour is the name of the late Prime Minister.

(2) On the basis of (1) I deduce that the late Prime Minister’s last name starts with a B.

(3) It turns out, unbeknownst to me, that the name of the late Prime Minister is Bannerman.

According to Russell, even though (2) is a true belief, it is not knowledge. Because (2) was inferred from a false belief, Russell would contend that it could not rightly be called knowledge. Knowledge, Russell argues, “is what is deduced from *known* premisses.”¹⁵⁴

With Russell’s analysis in place, we can address the actual Gettier cases:¹⁵⁵

Case 1:

(4) I am justified in believing that Jones will get the job, and that Jones has ten coins in his pocket.

(5) On the basis of (4) I deduce that the man with ten coins in his pocket will get the job.

¹⁵³ Russell 2004: 99.

¹⁵⁴ Russell 2004: 99.

¹⁵⁵ Gettier 1963.

(6) It turns out, unbeknownst to me, that I have ten coins in my pocket and I will get the job.

Case 2:

(7) I am justified in believing that Jones owns a Ford.

(8) On the basis of (7) I deduce that either Jones owns a Ford, or Brown is in Barcelona.

(9) It turns out, unbeknownst to me, that Brown is in Barcelona.

In each of the Gettier cases there is a justified belief – (4) and (7) – that entails another proposition – (5) and (8) respectively. The purported difficulty is that one can have a justified but false belief, which entails some other belief, which by sheer chance, turns out to be true. Gettier concludes that in both Case 1 and Case 2, one would have a JTB but not knowledge.

Clearly, though, given Russell's account of knowledge, neither case would be an instance of knowledge. Both (5) and (8) would be instances of true belief, but both also involve a move from a false belief to its logical entailment. To reiterate, "it is clear that a true belief is not knowledge when it is deduced from a false belief."¹⁵⁶

While the actual Gettier cases can be easily dismissed given the Russellian analysis, other Gettier-style cases that do not involve an explicit deduction from a false premise require further consideration. The challenge for a Russellian view is to retain some sense of an inference in which the premises are false, while not requiring that each inference be a conscious deduction from premises to a conclusion.

¹⁵⁶ Russell 2004: 99.

Fake-Barn County Case:¹⁵⁷

(10) I am justified in my belief that the thing I am pointing at is a very nice barn.

(11) Unbeknownst to me, I am driving in Fake-Barn County, where one in every four barns is real and the rest are barn-façades.

(12) As it turns out, the thing I am pointing at is a barn (in fact, a very nice one).

In this case, I do succeed in pointing to and picking out a very nice (real) barn, yet it is only a matter of luck that I pointed to an actual barn and not a barn-façade. This case requires no inference from (10) to some further conclusion; if (10) is justified, then (10) should be known (since it turned out to be true). However, the luck involved in the truth of (10) seems to count against its being an instance of knowledge.

Showing that (10) is based on false premises will show that (10) is not known, even though it turns out to be true. If the justification for (10) is one's perceptual evidence, then it would seem that one would also believe the following:

(13) There is no reason to doubt my perceptual experiences right now.

Or, in other words,

(14) I am in a normal kind of environment.

When one is in Fake-Barn County there is reason to doubt one's perceptual experiences, so both (13) and (14) would be false. In a sense, then, someone who believes (10) would do so on the basis of (13) and (14). But if this is true, then (10) cannot be called knowledge, since it depends on false premises ((13) and (14)).

Zagzebski Case:¹⁵⁸

¹⁵⁷ This case is a paraphrased version of the Gettier-style case that appears in Goldman 1976.

(15) Dr. Jones is justified in believing Smith has virus BAD.

(16) Unbeknownst to Dr. Jones, Smith's symptoms are the result of virus
WORSE.

(17) It turns out that Smith has contracted virus BAD, but so recently that he has
yet to show any BAD symptoms.

Since Dr. Jones believed Smith has virus BAD on the basis of very good
evidence, and since Smith did turn out to have the virus, it should be an instance of JTB.
But, as Zagzebski argues, "in this case, Dr. Jones' belief that Smith has virus [BAD] is
true, justified, and undefeated, but it is not knowledge."¹⁵⁹

Zagzebski's case would fail to be a counterexample if it can be shown that (15) is
not an instance of knowledge. The problems with (15) can be illustrated by the following
belief that Dr. Jones would also have, given the Zagzebski case:¹⁶⁰

(18) There are no other viruses that would cause the symptoms which are
plaguing Smith.

On Zagzebski's description of the case, (18) is false. But, if (18) is the reason that
Dr. Jones has for believing (15), then Dr. Jones does not know that Smith has virus BAD.
As a result, Dr. Jones has a belief that turns out to be true, but is not an instance of
knowledge, and consequently not a counterexample to the JTB account of knowledge.

¹⁵⁸ This case is a paraphrased version of the Gettier-style case that appears in: Zagzebski 2008.

¹⁵⁹ Zagzebski 2008: 211.

¹⁶⁰ In fact, in order to justify Dr. Jones's belief, Zagzebski asks (2008: 210) that we
"suppose that the symptoms are not compatible with any other known virus...."

It would be an understatement to say that the Russellian answer to the Gettier problem that I have sketched above is unpopular. The reason for this is that the account seems to over-intellectualize what is actually going on in our heads. In order for the Russellian response to work, one must have all sorts of dispositional beliefs – beliefs we actually have but are often unaware of – that are doing epistemic work.¹⁶¹

To see why the account of degenerate justification helps make sense of the Russellian response, one need only consider the weakened requirements that constitute degenerate justification. To be degenerately justified in a belief, one's current psychological state must cause – in the relevant way – the belief in question. Furthermore, as I have already argued, dispositional beliefs can partially constitute that cause. However, if one's belief is caused by an unjustified dispositional belief, then as Russell pointed out, the resulting belief does not constitute knowledge. In the Gettier cases discussed above, part of the relevant cause of one's belief is an unjustified dispositional belief, and consequently, the resulting belief also fails to be (even degenerately) known.

5.1.3 Epistemic Descent vs. Epistemic Ascent

Before showing why I take my account of epistemic descent to be preferable to the account of epistemic ascent Sosa develops, it is necessary to point out how the two views differ. The fundamental difference, which motivated what I called the apparatus of epistemic descent, is a difference in the architecture of our accounts.

Sosa's account, I argued, is best understood as being a bottom-up approach aimed at developing a philosophically satisfying account of knowledge (i.e., reflective

¹⁶¹ See McGrew and McGrew 2007: 7–34.

knowledge) from a less epistemically demanding account of knowledge (i.e., animal knowledge). In essence, Sosa builds from a less demanding account of knowledge to a more demanding one. My account is best understood as being top-down. I start with a philosophically satisfying account of justification (ideal justification) and then strip away epistemic layers in an effort to develop less epistemically demanding accounts of justification. My approach is to start with the more demanding account of knowledge and move down to one that is less demanding.

At this point Sosa could very well cry foul. What have I *added* to the discussion by simply restating what I granted was *his* distinction – the distinction between kinds of knowledge (or, as I have argued, justification)? Do the two views accomplish the same thing? Should someone who finds my approach compelling be just as compelled by his approach? I think the answer to each of these questions is “No.”

There is at least one other important difference that I want to highlight before showing why my account of epistemic descent is preferable to Sosa’s account of epistemic ascent. As I have already argued,¹⁶² Sosa commits himself to the idea that reflective knowledge just is “apt belief aptly noted.” On this conception of reflective knowledge one gets the better kind of knowledge by having apt belief about a lower-level belief. Thus, according to Sosa, one is able to move up levels through more apt belief. However, if this is right, then it follows that there is no such thing as the kind of ideal

¹⁶² See section 4.1.2.

knowledge that I start with. On Sosa's view, one can simply move up a potentially infinite number of levels by way of apt belief.¹⁶³

This conception of moving up levels, where an upper-limit conceptually fails to apply, is drastically different than the account I have offered. On my view, one starts with the ideal account of justification, where by 'ideal' I mean the conception of justification that cannot get (epistemically) any better. On my view, when one has a belief that is ideally justified, one has everything that one could want in terms of justified belief. Clearly, this is different from Sosa's account, which leaves room for an infinite series of increasingly (epistemically) better types of knowledge/justification.

I have already argued that Sosa's account of epistemic ascent is implausible.¹⁶⁴ I have also shown how the resources of my account of epistemic descent can be used to answer many of the challenges that threaten internalist accounts of justification. The account of epistemic descent that I have defended has the resources necessary to avoid the new evil demon problem, as well as the problem of easy knowledge.

First, consider the new evil demon problem. As I argued in Chapter 2, the uncontroversial assumption is that the justificatory status of the beliefs of an individual in a non-demon world would be the same as the justificatory status of the beliefs of an individual in a demon world, when the two individuals have had phenomenally indistinguishable lives. It is my contention that the account of ideal justification that I have defended is consistent with this uncontroversial assumption.

¹⁶³ Of course, my claim here is not that any human could actually move up an infinite number of levels, but instead that such a possibility is entailed by Sosa's conception of reflective knowledge.

¹⁶⁴ See my discussion of the new evil demon problem and the problem of easy knowledge.

Because the assurance-internalist account of ideal justification that I defend restricts the constituent elements of justification to features of immediate experience, two phenomenally indistinguishable individuals will have beliefs with the same justificatory status. Since justification, according to the assurance internalist, is ultimately parasitic upon the evidential connections that hold between psychological states and our beliefs, it follows that the justificatory status of a set of immediate experiences is consistent across possible worlds. As a result, it is irrelevant as far as ideal justification is concerned, whether an agent is in a demon world or not.

Perhaps most importantly, my account of degenerate justification will also accommodate the intuitions that motivate the new evil-demon problem. As I have already said, two individuals with phenomenally indistinguishable lives (“twins”) will have ideally justified beliefs with an identical status.¹⁶⁵ If one moves to a discussion of degenerate justification and compares two such individuals, one of whom is in a non-demon world while the other is in a demon world, it might turn out that the epistemic status of their degenerately justified beliefs is different. Since the condition of an appropriate causal role is an essential element in my account of degenerate justification, it follows that two individuals could have phenomenally indistinguishable lives but have beliefs that result from different (and sometimes deviant) causal chains. If this were the case, the individual whose beliefs are subject to deviant causal chains would not be degenerately justified, even if the twin (whose beliefs result from the appropriate causal chains) is degenerately justified in the same beliefs.

¹⁶⁵ Recall that ideal justification is an account of propositional justification.

Of course, it could be stipulated that both individuals have identical phenomenal lives and psychological states that result from identical causal processes. If this were the case, then the justificatory status of the individuals' degenerate justification would be identical. It is also important to note that this result is what one should expect. The point of the new evil-demon problem was not to draw attention to the difference between propositional and doxastic justification. Rather, as my account of degenerate justification makes clear, degenerate justification is ultimately dependent on, and parasitic upon, my conception of ideal justification. As I have pointed out, the evidential relations that do the epistemic work in both ideal and degenerate justification hold necessarily – across all possible worlds.

Next, consider the problem of easy knowledge. As I argued in Chapter 2, having too easy an answer to certain skeptical worries is itself worrisome. Using one's perceptual faculties to justify the reliability of one's perceptual faculties seems not uncontroversially problematic. But what, exactly, is wrong with using perception to justify perception, or memory to justify memory? Perhaps it is simply that, in establishing the efficacy of a belief-producing faculty, one should not appeal to the faculty in question.¹⁶⁶

While I think the kind of bootstrapping or circularity that the problem of easy knowledge makes explicit is generally problematic, it also seems to me that things are different when one considers the kind of ideal justification I defend and its reliance on acquaintance. In the case of acquaintance, and according to the account of noninferential

¹⁶⁶ Cohen 2002, for example, worries that all foundationalist accounts of justification face the problem of easy knowledge. However, the following section will show why this is not the case when one appeals to the appropriate account of acquaintance.

justification developed in Chapter 3, when one has a justified belief, one has before one's consciousness the very features that constitute not only justified belief, but also the sense of assurance that accompanies justified belief. It is important to remember that on the account I have defended, acquaintance is not itself a justification-conferring relation, but justification results when we are acquainted with the appropriate objects of immediate experience. When I am noninferentially justified in my belief that P, I am acquainted with the truth bearer (the thought that P), the truth maker (the fact that P), and the relation of fit obtaining between the truth bearer and the truth maker. As a result, I am able to *see* what my belief has going for it, and there is no further question as to whether or not my knowledge by acquaintance is working – everything is immediately before consciousness.

The use of acquaintance *with acquaintance* in justifying the efficacy of acquaintance is then importantly different from the use of perception to justify perception or memory to justify memory. For it is only in the case of acquaintance that one can have before one's consciousness all the epistemic elements necessary for a philosophically satisfying account of justification. Compare this result with, for example, a crude version of reliabilism. According to the reliabilist, when I am noninferentially justified in my belief that P, the belief that P need only be the result of a reliable belief-forming process. Of course, when I consider my belief that P, I might surely wonder whether the process that led to the belief was reliable.

5.2 Is this a Satisfying Response to Skepticism?

It should be clear by now that I think skeptical challenges offer a serious threat to the limits of ideal justification. We simply do not have before our consciousness the

resources to justify many of the beliefs we commonly take to be justified. Even if we did have before consciousness everything that was required to justify those beliefs, the required inferences would be so varied and so complex that it seems likely that many of us are simply not up to the cognitive challenge.

Fortunately, all is not lost. As we saw in Chapter 1, Goldman argued that because skepticism is the inevitable result of internalism, internalism should be rejected. However, the account of degenerate justification I have developed in the previous chapters shows why Goldman's skeptical worries fail. The internalist need not conclude that she must give up internalism (or for that matter, adopt skepticism). Perhaps the most radical conclusion one should adopt on the basis of my view is agnosticism with regard to many of our commonsense beliefs.

Still, it is important that the success of this project is cast in the appropriate light. One may have been hopeful at the outset of this project that I was going to deliver an account of ideal justification that brought with it a guarantee of truth, and that one could use that sense of ideal justification to *see* that one's degenerately justified commonsensical beliefs really were justified. However, this is not my project, and certainly not my conclusion.

Since it is possible for even an ideally justified belief to be false, it simply follows that (since knowledge requires truth) it is possible that one could have ideal justification and still not have knowledge.¹⁶⁷ Because of the requirements on ideal justification, we are simply not in a position to know (or even to be ideally justified in believing) that we have

¹⁶⁷ It does not follow from this that we can never know that we know something.

degenerate justification. We are simply not in a position to see that the appropriate causal relations are obtaining.

Perhaps all I can say in the case of degenerate justification is that if our beliefs are caused in the appropriate way, then we have degenerately justified belief. Thus, it *might* be the case that many of our beliefs are degenerately justified. Of course, saying something *might* be the case does not necessarily mean it is the case (or that it is even likely). As a result, the skeptic who concedes the possibility of degenerate justification might still argue that the likelihood that such an account of justification is true is extremely low. One might worry at this point that I have served up a version of externalism in a sheep's clothing.¹⁶⁸

It seems that many of the beliefs we commonly take to be justified could very well be justified, albeit degenerately. At this point, one should ask, *What reason do we have to think that the appropriate causal chains required for degenerate justification obtain?* Here, the appropriate answer depends on the audience. It would surely be a bold question coming from a paradigm reliabilist externalist. The account of degenerate justification I have offered is at least as well-suited to answer skepticism as the reliabilist is.¹⁶⁹ If the question comes from a committed skeptic, I am afraid there might not be a satisfying response. I have not shown how the most philosophically interesting conception of justification – the ideal justification offered by assurance internalism – can justify the vast majority of our commonsensically justified beliefs. Nor have I shown that

¹⁶⁸ I do not think this conclusion is correct. See section 4.2.3 for an account of why I think degenerate justification is still appropriately labeled an account of internalist justification.

¹⁶⁹ In fact, as I argued in the previous section, I think my account is better suited to respond to the skeptic than a paradigm reliabilist.

the vast majority of our commonsensically justified beliefs do in fact enjoy degenerate justification. Of course, this was not my goal.

My accounts of degenerate justification are parasitic upon (i.e., derived from) ideal justification. As a result, if the accounts of ideal justification that I defend are not possible, then my accounts of degenerate justification lose their grounding. The successful defense of ideal justification plays a crucial role in a full presentation of my account of degenerate justification. One issue that was set aside earlier¹⁷⁰ had to do with the role of the principle of inferential justification in my account of ideally justified inferential belief. In order to have an ideally justified inferential belief that P on the basis of E, one must have the justified belief that E and the justified belief that E makes probable P.

I argued that one had better be able to make sense of having a justified belief that E makes probable P in way that does not invite a vicious regress. On my account, the appropriate strategy for trying to make sense of how we can have a justified belief between one's evidence and one's conclusion required being able to see the evidential connection that holds between E and P. In addition, I argued that these evidential connections are necessary *a priori* truths – they are true in all possible worlds.

I have also argued that the necessity of these evidential connections is one of the most advantageous parts of my view. It is the necessity of these evidential connections that allows me to answer the new evil-demon problem – a problem that I argued Sosa's view cannot handle – and thus helps shows why my apparatus of epistemic descent is preferable to Sosa's account of epistemic ascent.

¹⁷⁰ See section 4.2.2.

Justified belief in these evidential connections, necessary truths, and relations between premises and conclusion are of utmost importance to the account I have offered. This justified belief must be noninferential, lest we fall into the clutches of a vicious regress. How might we attempt to explicate justified belief of these evidential relations? This question, perhaps not surprisingly, will depend a great deal on the metaepistemological presuppositions of the philosopher defending an account of justified belief in the required evidential relations. Clearly, depending on the epistemic principles¹⁷¹ to which one is willing to commit oneself, the nature and scope of the defense will vary widely.

For instance, one might only be willing to commit oneself to the epistemic principles that govern deductive inference. For example, it seems that I can know without inference the rule of *simplification* – from P and Q, one can infer P. Perhaps the required evidential relation can be known in the case of deductive inference, where I am able to see the appropriate evidential connection. Consider again the simple case of ideally justified inferential belief from section 4.2.2:

I have the noninferentially justified belief that P. I have the justified belief that P entails that either P or Q (because I have a noninferentially justified belief in the rule of *addition* – from P one can infer P or Q). As a result I have the ideally justified inferential belief that either P or Q.

If we restrict ideally justified inferential belief to only those inferences that involve a simple deductive relation between premises and conclusion, the extent of ideally justified inferential belief will be severely limited. I am willing to grant as much, however. Even if this is the case, I have already succeeded in showing that ideal inferential justification is

¹⁷¹ I am thinking of epistemic principles as specifying conditions under which it is rational to form belief.

possible, even if it is not very common. Since I have shown that it is possible, the conceptions of degenerate justification that I defend are possible as well.

Some may not be happy with the conclusion that we can only have ideally justified inferential beliefs in cases of simple deductive inferences. One might very well try to motivate other epistemic principles in an effort to greatly increase the kinds of justified inferences that can be used in relating premises to a conclusion.

One might be interested in trying to expand one's stock of the epistemic principles that allow us to ground justified belief in the evidential connections used in ideally justified inferential belief. This could be done by moving from the kinds of simple deductive inferences discussed above to inductive inferences. One might, following Hume, consider the rationality of inductive inferences that involve constant (or near-constant) conjunction. We see smoke and think fire; we hear the growing sound of footsteps and think someone is coming towards us. Perhaps, then, it is my past experience of constant conjunction between a certain kind of experience and a certain state of affairs.

However, as Hume was well aware, it would seem that our knowledge of the epistemic principle of induction is known inferentially, and this is a problem if we hope to avoid a regress. The reason we infer from the existence of one thing to that of another, argued Hume, is because of our experience of constant conjunction in the past. Thus, smoke (almost) always came with fire in the past, the growing sound of footsteps (almost) always came with a person causing them. But how can we justify the use of constant conjunction, without appealing to past experience?

Even if Hume has shown that constant conjunction will not get the job done, there are plenty of other epistemic principles that might be advanced in an effort to ground

inductive inference. One could, for instance, argue for an epistemic principle along the lines of phenomenal conservatism:

PC: If it seems to S that P, then S is justified in inferring that P.

On this construal, the principle PC is clearly inferential.¹⁷² Thus, if I have the justified belief that I seem to be at the beach, then according to PC I am justified in inferring that I am at the beach. Clearly, if one were able to get PC off the ground, it would provide a very powerful response to the skeptic.

Advocates of other varieties of phenomenal conservatism, like Huemer, argue that one should treat seemings consistently: all seemings – including but not limited to perceptual and memorial seemings – are legitimate sources of justification that can be used with a principle like PC. If one were to adopt Huemer’s “compassionate” approach to seemings, and one were to adopt the above formulation of PC, it is clear that one could have ideally justified inferential belief in many of the cases in which we commonly take ourselves to be justified.

Although such a formulation of PC could be a useful tool against the skeptic, it is still far from clear how one might be noninferentially justified in the belief that PC is a legitimate form of inference. It is just not clear to me why one would think that it is necessary that seemings make likely true belief. Why it is that just because something seems a certain way, we should be inclined to think it is in fact that way? The advocate of PC would need to offer an analysis of seemings such that the seeming that P makes probable P.

¹⁷² This is not the kind of principle someone like Huemer would advance. Huemer’s account of phenomenal conservatism is noninferential: “If it seems to S that p, then, in the absence of defeaters, S thereby has at least some degree of justification for believing that p” 2007: 30.

An alternative approach to grounding inductive inference might rely on an Keynesian account of probability. One might argue that probability relations are best understood as relations that hold between propositions necessarily. On this interpretation of probability, it might be argued that deductive inferences are just the upper bound of the making-probable relation. Thus, in the very same way that one is able to see deductive entailments (or to be noninferentially justified in believing that they obtain), one is able to see probability relations. Just as was the case with our simple deductive inference, if one is noninferentially justified in believing that P, and one is able to see the making-probable relation that stands between P and some further belief Q, one is inferentially justified in the belief that Q. On this view, however, one is required to see the objective probability relations, and whether or not we have such an ability is controversial to say the least.

Another approach intended to account for inductive inference relies on inference to the best explanation (IBE). It could be argued that explanatory notions of simplicity, beauty, or consilience can be adequately tied to the making-probable relation. Thus, one might argue that because the considerations that pick out the best explanation are also tied to explanations that are likely to be true, we can infer from the fact that an explanation is the best to the conclusion that it is probably true. Of course, in order for this type of approach to work, the advocate of IBE would need to noninferentially justifiably believe that the explanatory virtues are truth-conducive. Indeed, to avoid an infinite regress, the IBE advocate would have to know they are necessarily truth-conducive.

Finally, one might try to use an objectivist account of probability to reconstruct *abductive* inferences. The IBE model can be elucidated by appeal to Peirce's formulation of abduction:

The surprising fact E is observed;
 But if H were true, E would be a matter of course,
 Hence, there is reason to suspect that H is true.¹⁷³

By adding a further step to Peirce's schema, we can arrive at the more robust notion of IBE: in the case where there is a set of possible hypotheses that explain the evidence, these hypotheses must be compared. Thus, in an instance of IBE, it is granted that some evidence has been observed, and that evidence needs an explanation (the surprising fact E is observed). All of the possible explanations would make the occurrence of E predictable, but if H were true, E would be a matter of course. It is inferred that there is some evidence for each hypothesis – there is reason to suspect that H is true. Finally, of the possible explanations, one is preferable; it is then concluded that, probably, the best explanation is the right one.

Despite the intuitive appeal of IBE, there are two difficulties that raise significant doubts as to the legitimacy of this form of inference. First, it could be argued that the best explanation might be only the best of a bad lot. Second, one may wonder why it is that the criteria used for picking an explanation should be truth-conducive at all.

I have given an account of ideal justification that shows why at least some of our beliefs are ideally justified.¹⁷⁴ In addition, the account I have offered seems to carve out

¹⁷³ This formulation of Peirce's schema is presented in McGrew 2003. See Peirce 1935: 5.189. for the original formulation.

an appropriate role for the philosopher. It is, after all, not the job of the philosopher (*qua* philosopher) to figure out what an individual *actually* knows. Instead, the job of the philosopher (*qua* philosopher) is to figure out what *would* count as an instance of knowledge – to figure out what the necessary and sufficient conditions of knowledge are. This is the task that I have undertaken. I have sketched an internalist account of justification that does not commit the internalist to skepticism, yet it is also an account that takes skeptical challenges seriously.¹⁷⁵

¹⁷⁴ As I argued previously, my account has distinct advantages over a rival account like Sosa's because I have shown that at least some of our beliefs are ideally justified and not susceptible to new evil demon type scenarios.

¹⁷⁵ I would like to thank my advisor, Richard Fumerton, for his patience and guidance as I worked through this project. Without his help what I have said would be even harder to understand.

REFERENCES

- Almeder, Robert. 1995. "Externalism and Justification." *Philosophia* 24. 465–69.
- Alston, William. 1983. "What's Wrong with Immediate Knowledge?" *Synthese* 55. 73–95.
- Alston, William. 1986. "Internalism and Externalism in Epistemology." *Philosophical Topics* 14. 179–221.
- Alston, William. 1988. "An Internalist Externalism." *Synthese* 74. 265–83.
- Alston, William. 1991. *Perceiving God*. Ithaca: Cornell University Press.
- Alston, William. 1993. *The Reliability of Sense Perception*. Ithaca: Cornell University Press.
- Alston, William. 1995. "How to Think About Reliability." *Philosophical Topics* 23. 1–29.
- Alston, William. 2004. "The Challenge of Externalism." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 37–52.
- Alston, William. 2005. *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Armstrong, David. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Audi, Robert. 1993. *The Structure of Justification*. Cambridge: Cambridge University Press.
- Audi, Robert. 2003. *Epistemology*. Second Edition. New York: Routledge.
- Austin, J. L. 1961. *Philosophical Papers*. Oxford: Oxford University Press.
- Ayer, A. J. 1940. *The Foundation of Empirical Knowledge*. London: MacMillan.
- Ayer, A. J. 1956. *The Problem of Knowledge*. London: Penguin.
- Bergmann, Michael. 2000. "Deontology and Defeat." *Philosophy and Phenomenological Research* 60. 87–102.
- Bergmann, Michael. 2006. *Justification without Awareness*. Oxford: Oxford University Press.

- BonJour, Laurence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5. 3–73.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- BonJour, Laurence. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- BonJour, Laurence. 2002. *Epistemology*. Lanham, MD: Rowman & Littlefield.
- BonJour, Laurence and Ernest Sosa. 2003. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Hoboken, NJ: Blackwell.
- Butchvarov, Panayot. 1970. *The Concept of Knowledge*. Chicago: Northwestern University Press.
- Carroll, Lewis. 1895. "What the Tortoise Said to Achilles." *Mind* 4 (14). 278–80.
- Chisholm, Roderick. 1942. "The Problem of the Speckled Hen." *Mind* 51 (204). 368–73.
- Chisholm, Roderick. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Chisholm, Roderick. 1976. *Theory of Knowledge*. Second Edition. Upper Saddle River, NJ: Prentice Hall.
- Chisholm, Roderick. 1988. "The Indispensability of Internal Justification." *Synthese* 74. 285–96.
- Cohen, Stewart. 1984. "Justification and Truth." *Philosophical Studies* 46. 279–95.
- Cohen, Stewart. 1998. "Fumerton on Metaepistemology and Skepticism." *Philosophy and Phenomenological Research* 58 (4). 913–8.
- Cohen, Stewart. 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2). 309–29.
- Cohen, Stewart. 2009. "Knowledge as Aptness." *Philosophical Studies* 144. 121–5.
- Comesaña, Juan. 2002. "The Diagonal and the Demon." *Philosophical Studies* 110 (3). 249–66.
- Conee, Earl. 2004. "Externalism, Internalism, and Skepticism." *Philosophical Issues* 14. 78–90.

- Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Corbí, Josep. 2000. "The Principle of Inferential Justification, Scepticism and Causal Belief." Ernest Sosa and Enrique Villanueva, eds. *Skepticism*. Hoboken, NJ: Blackwell. 377–85.
- Cottingham, John, Robert Stoothoff, and Dugald Murdoch. 1986–91. *The Philosophical Writings of Descartes*. Cambridge: Cambridge University Press. Three volumes.
- Cruz, Joseph and John Pollock. 2004. "The Chimerical Appeal of Epistemic Externalism." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 125–42.
- Dancy, Jonathan. 1992. "Externalism for Internalists." *Philosophical Issues* 2. 93–114.
- Dretske, Fred. 1969. *Seeing and Knowing*. Chicago: University of Chicago Press.
- Dretske, Fred. 1971. "Conclusive Reasons." *Australasian Journal of Philosophy* 49. 1–22.
- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Fales, Evan. 1996. *A Defense of the Given*. Lanham, MD: Rowman & Littlefield.
- Fales, Evan. 2004. "Proper Basicity." *Philosophy and Phenomenological Research* 68. 373–83.
- Fales, Evan and Richard Fumerton. "Levels of Epistemic Justification." Unpublished.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, NJ: Pearson Education Inc.
- Feldman, Richard. 2004. "In Search of Internalism and Externalism." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 143–57.
- Feldman, Richard. 2005a. "Justification is Internal." Matthias Steup and Ernest Sosa, eds. *Contemporary Debates in Epistemology*. Hoboken, NJ: Blackwell. 270–84.
- Feldman, Richard. 2005b. "Respecting the Evidence." *Philosophical Perspectives* 19. 95–119.
- Feldman, Richard. 2006. "BonJour and Sosa on Internalism, Externalism, and Basic Beliefs." *Philosophical Studies* 131. 713–28.
- Fogelin, Robert. 1967. *Evidence and Meaning*. New York: Routledge.

- Foley, Richard. 1993. "What Am I to Believe?" Steven Wagner and Richard Warne, eds. *Naturalism: A Critical Appraisal*. South Bend, IN: University of Notre Dame Press. 147–62.
- Fumerton, Richard. 1976. "Inferential Justification and Empiricism." *The Journal of Philosophy* 73. 557–69.
- Fumerton, Richard. 1985. *Metaphysical and Epistemological Problems of Perception*. Lincoln: University of Nebraska Press.
- Fumerton, Richard. 1988. "The Internalism/Externalism Controversy." *Philosophical Perspectives* 2. 443–59.
- Fumerton, Richard. 1995a. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Fumerton, Richard. 1995b. "Sosa's Epistemology." *Philosophical Issues* 5. 15–27.
- Fumerton, Richard. 2002. "Exemplarizing and Self-Presenting States." *Philosophy and Phenomenological Research* 64 (2). 431–5.
- Fumerton, Richard. 2004. "Inferential Internalism and the Presuppositions of Skeptical Arguments." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 157–68.
- Fumerton, Richard. 2006a. *Epistemology*. Hoboken, NJ: Blackwell.
- Fumerton, Richard. 2006b. "Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament." Thomas Crisp, ed. *Knowledge and Reality: Essays in Honor of Alvin Plantinga*. Berlin: Springer. 179–91.
- Fumerton, Richard. 2011. "Reflective Knowledge and Intellectual Assurance." *International Journal for the Study of Skepticism* 1 (2). 113–23.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23. 121–3.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. Holland: Reidel.
- Goldman, Alvin. 1967. "A Causal Theory of Knowing." *The Journal of Philosophy* 64. 357–72.
- Goldman, Alvin. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73. 771–791.

- Goldman, Alvin. 1979. "What is Justified Belief?" George Pappas, ed. *Justification and Knowledge*. Berlin: Springer. 1–23.
- Goldman, Alvin. 1980. "The Internalist Conception of Justification." *Midwest Studies in Philosophy* 5. 27–51.
- Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, Alvin. 1988. "Strong and Weak Justification." *Philosophical Perspectives* 2. 51–69.
- Goldman, Alvin. 1993. "Epistemic Folkways and Scientific Epistemology." *Philosophical Issues* 3. 271–85.
- Goldman, Alvin. 1999. "Internalism Exposed." *The Journal of Philosophy* 96. 271–93.
- Goldman, Alvin. 2009. "Internalism, Externalism, and the Architecture of Justification." *The Journal of Philosophy* 106. 309–38.
- Greco, John. 1990. "Internalism and Epistemically Responsible Belief." *Synthese* 85. 245–77.
- Greco, John. 1993. "Virtues and Vices of Virtue Epistemology." *Canadian Journal of Philosophy* 23 (3). 413–432.
- Greco, John. 1999. "Agent Reliabilism." *Philosophical Perspectives* 13. 273–96.
- Greco, John. 2000. "Scepticism and Epistemic Kinds." Ernest Sosa and Enrique Villanueva, eds. *Skepticism*. Hoboken, NJ: Blackwell. 366–76.
- Greco, John. 2003. "Virtue and Luck, Epistemic and Otherwise." *Metaphilosophy* Vol. 34 No. 3. 353–365.
- Greco, John. 2004. "Externalism and Skepticism." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 53–64.
- Greco, John. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Grundmann, Thomas. 2004. "Counterexamples to Epistemic Externalism Revisited." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 65–76.
- Haack, Susan. 1993. *Evidence and Inquiry*. Hoboken, NJ: Blackwell.

- Harper, William. 1988. "Papier Mâché Problems in Epistemology: A Defense of Strong Internalism." *Synthese* 116. 27–49.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Henderson, David and Terry Horgan. 2001. "Practicing Safe Epistemology." *Philosophical Studies* 102. 227–58.
- Henderson, David and Terry Horgan. 2007. "Some Ins and Outs of Transglobal Reliabilism." Sanford Goldberg, ed. *Internalism and Externalism in Semantics and Epistemology*. Oxford: Oxford University Press. 100–30.
- Hookway, Christopher. 2000. "Scepticism and the Principle of Inferential Justification." Ernest Sosa and Enrique Villanueva, eds. *Skepticism*. Hoboken, NJ: Blackwell. 344–65.
- Huemer, Michael. 1999. "The Problem of Memory Knowledge." *Pacific Philosophical Quarterly* 80. 346–57.
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
- Huemer, Michael. 2002. "Fumerton's Principle of Inferential Justification." *Journal of Philosophical Research* 27. 330–40.
- Huemer, Michael. 2006. "Phenomenal Conservatism and the Internalist Intuition." *American Philosophical Quarterly* 43 (2). 147–58.
- Huemer, Michael. 2007. "Compassionate Phenomenal Conservatism." *Philosophy and Phenomenological Research* 74 (1). 30–55.
- Huemer, Michael. 2011. "The Puzzle of Metacoherence." *Philosophy and Phenomenological Research* 82 (1). 1–21.
- Hume, David. 1998. Tom Beauchamp, ed. *An Enquiry Concerning the Principles of Morals*. Oxford: Oxford University Press.
- Hume, David. 2000. David Norton and Mary Norton, eds. *A Treatise of Human Nature*. Oxford: Oxford University Press.
- Kornblith, Hilary. 1988. "How Internal Can You Get?" *Synthese* 74. 313–27.
- Kornblith, Hilary. 2001. *Epistemology: Internalism and Externalism*. Hoboken, NJ: Blackwell.

- Kornblith, Hilary. 2004. "Conditions on Cognitive Sanity and the Death of Internalism." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 77–88.
- Kyberg, Henry. 1983. *Epistemology and Inference*. Minneapolis: University of Minnesota Press.
- Landesman, Charles. 2002. *Skepticism*. Hoboken, NJ: Blackwell.
- Lehrer, Keith. 1974. *Knowledge*. Oxford: Oxford University Press.
- Lehrer, Keith and Stewart Cohen. 1983. "Justification, Truth, and Coherence." *Synthese* 55. 191–207.
- Lehrer, Keith. 2000. *Theory of Knowledge*, Second Edition. Boulder: Westview Press.
- Lehrer, Keith and David Truncellito. 2004. "Knowledge, Justification, and the Cooperative World." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 169–80.
- Leite, Adam. 2005. "Epistemological Externalism and the Project of Traditional Epistemology." *Philosophy and Phenomenological Research* 70. 505–33.
- Lewis, C. I. 1946. *An Analysis of Knowledge and Valuation*. Chicago: Open Court.
- Lycan, William. 1977. "Evidence One Does Not Possess." *Australasian Journal of Philosophy* 55. 114–26.
- McGrew, Timothy. 1995. *The Foundations of Knowledge*. Lanham, MD: Rowman & Littlefield.
- McGrew, Timothy and Lydia McGrew. 2007. *Internalism and Epistemology: The Architecture of Reason*. New York: Routledge.
- McLaughlin, Brian. 2000. "Self-Knowledge, Externalism, and Skepticism." *Proceedings of the Aristotelian Society* 74. 93–118.
- Moser, Paul. 1989. *Knowledge and Evidence*. Cambridge: Cambridge University Press.
- Moya, Carlos and Tobies Grimaltos. 2000. "Memory and Justification: Hookway and Fumerton on Scepticism." Ernest Sosa and Enrique Villanueva, eds. *Skepticism*. Hoboken, NJ: Blackwell. 386–94.
- Plantinga, Alvin. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Pollock, John. 1974. *Knowledge and Justification*. Princeton: Princeton University Press.

- Pollock, John and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. Lanham, MD: Rowman & Littlefield.
- Pritchard, H.A. 1950. *Knowledge and Perception*. Oxford: Oxford University Press.
- Reid, Thomas. 1997. *An Inquiry into the Human Mind on the Principles of Common Sense*. Edited by Derek Brookes. University Park, PA: Pennsylvania State University Press.
- Russell, Bertrand. 1914. *Our Knowledge of the External World: As a Field for Scientific Method in Philosophy*. Crows Nest, New South Wales: George Allen and Unwin.
- Russell, Bertrand. 1927. *Human Knowledge: Its Scope and Limits*. Crows Nest, New South Wales: George Allen and Unwin.
- Russell, Bertrand. 2004 (first printing 1912). *The Problems of Philosophy*. New York: Barnes and Noble.
- Selvin, Steve. 1975. "A problem in probability (letter to the editor)." *American Statistician* 29 (1). 67.
- Schantz, Richard. 2004. "Empiricism Externalized." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 89–112.
- Sosa, Ernest. 1988. "Beyond Skepticism, to the Best of our Knowledge." *Mind* 97 (386). 153–88.
- Sosa, Ernest. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Sosa, Ernest. 2004a. "Circularity and Epistemic Polarity." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 113–22.
- Sosa, Ernest. 2004b. "Replies." John Greco, ed. *Ernest Sosa: and His Critics*. Malden, MA: Blackwell. 275–326.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. New York: Oxford University Press.
- Sosa, Ernest. 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. New York: Oxford University Press.
- Sosa, Ernest. 2011. *Knowing Full Well*. Princeton: Princeton University Press.

- Steup, Matthias. 1999. "A Defense of Internalism." Louis Pojman, ed. *The Theory of Knowledge: Classical and Contemporary Readings*, Second Edition. Stamford, CT: Wadsworth. 373–84.
- Steup, Matthias. 2000. "Doxastic Voluntarism and Epistemic Deontology." *Acta Analytica* 15. 25–56.
- Steup, Matthias. 2004. "Internalist Reliabilism." *Philosophical Issues* 14. 403–25.
- Stroud, Barry. 1994. "Scepticism, 'Externalism', and the Goal of Epistemology." *Proceedings of the Aristotelian Society* 68. 291–307.
- Stroud, Barry. 2004. "The Epistemological Promise of Externalism." Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 181–92.
- Sutton, Jonathan. 2007. *Without Justification*. Cambridge, MA: MIT Press.
- Williams, Michael. 2001. *Problems of Knowledge*. Oxford: Oxford University Press.
- Williams, Michael. 2004. "Is Knowledge a Natural Phenomenon?" Richard Schantz, ed. *The Externalist Challenge*. New York: Walter de Gruyter. 193–209.
- Zagzebski, Linda. 2008. "The Inescapability of Gettier Problems." Ernest Sosa, et al., eds. *Epistemology: An Anthology*. Second Edition. Malden, MA: Blackwell.