
Theses and Dissertations

Summer 2011

The skeptic's dogmatism: a constructive response to the skeptical problem

Kaplan Levent Hasanoglu
University of Iowa

Copyright 2011 Kaplan Levent Hasanoglu

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/1144>

Recommended Citation

Hasanoglu, Kaplan Levent. "The skeptic's dogmatism: a constructive response to the skeptical problem." PhD (Doctor of Philosophy) thesis, University of Iowa, 2011.
<https://ir.uiowa.edu/etd/1144>.

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



Part of the [Philosophy Commons](#)

THE SKEPTIC'S DOGMATISM: A CONSTRUCTIVE RESPONSE TO THE
SKEPTICAL PROBLEM

by
Kaplan Levent Hasanoglu

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 2011

Thesis Supervisor: Associate Professor Evan Fales

ABSTRACT

The problem of philosophical skepticism relates to the difficulty involved in underwriting the claim that we know anything of spatio-temporal reality. It is often claimed, in fact, that proper philosophical scrutiny reveals quite the opposite from what common sense suggests. Knowledge of external reality is thought to be even quite obviously denied to us as a result of the alleged fact that we all fail to know that certain skeptical scenarios do not obtain. A skeptical scenario is one in which we have neurological occurrences just like any normal situation in which we actually perceive spatio-temporal objects, but where we are deceived in some sense as a result of the manipulation of our brains from some outside source. In this work I attempt to address the problem of philosophical skepticism by claiming that most of us are able to come to know plenty about external reality, since we can come to realize that a certain philosophical theory of perception is correct. The theory I have in mind is what I call a non-cognitive theory of perception. According to this view, perceptual experience is defined by continuous, sensitive behavioral interaction with spatio-temporal objects of the appropriate size, shape, hardness, speed, etc. Knowing that this theory of perception is correct is equivalent to knowing that we know plenty about external reality. This is ultimately because by knowing that a non-cognitive theory of perception is true, we know that any skeptical scenario must fail to obtain.

The structure of the work proceeds by first discussing the significance of the problem of philosophical skepticism in some detail. Chapter 1 lays out how the problem does indeed forcefully arise if it is conceded that we fail to know that the skeptical scenarios fail to obtain. Chapter 2 develops the sort of view of our epistemological

situation that falls out of accepting without qualification that the problem exists. Chapter 3 examines and criticizes certain popular responses to the skeptical problem. The main goal of these first three preliminary chapters is to indicate that once it is admitted that we fail to know that the skeptical scenarios fail to obtain, the problem of philosophical skepticism forcefully presents itself.

Chapter 4, however, attacks the idea that we fail to know that the skeptical scenarios fail to obtain. In this chapter I argue that this idea is wedded to the position that a certain theory of perception is correct. The theory in question is what I call a conjunctive theory of veridical experience. According to this view, normal experience of spatio-temporal objects occurs when a subject has a certain perceptual experience, and that experience also happens to match up with what is really the case. Only when a theory of this sort is assumed, I argue, is the claim that we fail to know that the skeptical scenarios fail to obtain found to be obvious. In chapter 5 I argue that, in fact, a non-cognitive theory, rather than a conjunctive theory, is the correct view to maintain. In the final chapter 6 I develop the epistemological position that falls out of accepting a non-cognitive theory.

Abstract Approved: _____

Thesis Supervisor

Title and Department

Date

THE SKEPTIC'S DOGMATISM: A CONSTRUCTIVE RESPONSE TO THE
SKEPTICAL PROBLEM

by

Kaplan Levent Hasanoglu

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 2011

Thesis Supervisor: Associate Professor Evan Fales

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

Kaplan Levent Hasanoglu

has been approved by the Examining Committee
for the thesis requirement for the Doctor of Philosophy
degree in Philosophy at the July 2011 graduation.

Thesis Committee: _____
Evan Fales, Thesis Supervisor

Richard Fumerton

Gregory Landini

Ali Hasan

Eliot Hazeltine

TABLE OF CONTENTS

CHAPTER

I.	THE PROBLEM OF PHILOSOPHICAL SKEPTICISM.....	1
	1.1 Introduction.....	1
	1.2 The Dream Argument.....	7
	1.3 Fallibilism Developed.....	10
	1.4 Pessimistic Fallibilism Conditionally Defended.....	29
	1.5 Wright's Argument Against Radical Skepticism.....	48
	1.6 Conclusion.....	58
II.	PESSIMISTIC FALLIBILISM DEVELOPED.....	60
	2.1 Introduction.....	60
	2.2 Assertion and Negative Epistemological Arguments.....	63
	2.3 Conclusion.....	85
III.	NON-DISCURSIVE OPTIMISM.....	87
	3.1 Introduction.....	87
	3.2 Wittgenstein.....	99
	3.3 Contextualism.....	109
	3.4 Process-Reliabilism.....	123
	3.5 Epistemological Naturalism.....	146
	3.6 Conclusion.....	158
IV.	DOUBTS ABOUT FALLIBILISM.....	160
	4.1 Introduction.....	160
	4.2 Cartesian Questions.....	162
	4.3 Moore's Internal Answer to Skepticism.....	171
	4.4 Against the Conditional Correctness of Skepticism.....	173
	4.5 The Problem of Philosophical Skepticism.....	184
V.	SECOND-ORDER SKEPTICISM REFUTED.....	190
	5.1 Introduction.....	190
	5.2 A Non-Cognitive Theory of Perception.....	190
	5.3 The Main Argument.....	194
	5.4 Conclusion.....	208
VI.	INFALLIBILISM DEFENDED.....	210
	6.1 Introduction.....	210
	6.2 Fallibilism, Internalism and Externalism.....	211
	6.3 Infallibilism Defended.....	223

NOTES.....	239
REFERENCES	249

CHAPTER I THE PROBLEM OF PHILOSOPHICAL SKEPTICISM

1.1 Introduction

If perception cannot give us knowledge of the contingent makeup of external reality, then nothing can.¹ The philosophical project of determining the extent of our external world knowledge involves nothing more than determining what knowledge of this type we can squeeze out of our perceptual experiences. By reasonably answering certain questions about perceptual experience, one thereby makes just as reasonable certain fundamental epistemological beliefs.² In this sense, metaphysics informs epistemology; but this is not to say that it *has priority* over epistemology. Merely being willing to assert a certain view of perception (at least in a philosophical context) implies that one finds it most reasonable to maintain that position.³ So, reciprocally, in this sense (along with a more familiar one) epistemology informs metaphysics as well.⁴ The right way to think about this matter seems to be that neither epistemology nor metaphysics takes priority over the other.

This logical interdependence between metaphysics and epistemology can be turned against the radical external world skeptic. Here's how. Radical skepticism--the position that we fail to justifiably believe anything about the external world--is often championed as the view that inevitably results from a properly philosophical examination of the epistemological status of our external world beliefs.⁵ Consequently, it is in the interest of one who wants to grapple with the skeptic to force the latter to get clear on just what it means to perform this sort of examination. A failure to be clear on this matter would obviously infect the clarity of the view that skepticism is the most reasonable philosophical position. More strongly, if the skeptic can't demonstrate how his position

follows *for clear and demonstrable reasons*, then he can easily come out sounding like the crazy one in the room. Psychologically speaking, epistemological realism regarding the mind-independent external world--which is the position that we know plenty about external reality--is all but forced on us. Skepticism, however, is not a view about what we are/are not naturally psychologically compelled to think regarding what we know. Or, at least, it is more complicated than that.

In any case, I hope to show in what follows that after we *do* get clear on the precise rules of the game the skeptic wishes to play, we will see that skepticism *can't* be defended. In fact, the proper basis for the denial of skepticism is even anticipated, in a way, by the still prevailing contemporary view that there is something *paradoxical* about the skeptical result.⁶ Paradoxes pack a normative punch; their very existence cries out for a smoothing out, one way or the other, of the cognitive dissonance. In the case of the skeptical result the paradox results, in part, from the denial of common sense knowledge claims. Therefore, by making good on the negative claim that no good reasons exist for denying these claims, one makes the way to resolve the otherwise paradoxical matter rather clear.

It is the (eventual) burden of this work to make the moving parts of this kind of resolution obvious. In chapters 4 - 6 I will argue that in the context of the philosophical examination of the extent of our external world knowledge, the skeptical thesis can be found to be just as absurd as what I will call a *non-cognitive theory of perception* is found plausible. By this sort of theory I mean, first of all, something negative. A non-cognitive theory denies, that is, what I call a *conjunctive theory of veridical experience*. According to a conjunctive theory, veridical experience is logically complex in the following sense:

in order to veridically perceive it must *both* be true that one is having certain sense experiences *as well as* that those experiences "match up with" or "are satisfied by" what is really the case.⁷ Rather than taking perceptual experience to be logically complex in this (obscure) sense, a non-cognitive theorist takes it to be constituted by the exercising of sense-organs (eyes, ears, etc.) so as to allow one to comport one's body appropriately given one's surrounding environment.⁸

In later chapters, I'll have much more to say by way of clarification of this theory. For now it is only important to note how arguing for a non-cognitive theory in the manner I intend makes use of the fact that epistemology informs metaphysics (in the sense discussed above): a non-cognitive theory is, I hope to show, far and away the most reasonable position to hold in the context of a philosophical assessment of the extent of our external world knowledge. The fact that metaphysics informs epistemology will also be pressed into service: adopting a non-cognitive theory, I argue, provides one with the resources for maintaining a form of *infallibilism*. Infallibilists of this sort claim that human beings are able, at least in certain cases, to conclusively verify contingent external world beliefs. They claim that we are sometimes able to verify--via perceptual experiences and perceptually-regulated acts--that all cases where matters are deceptively otherwise have not obtained. Briefly put, a non-cognitive theory supports infallibilism by, most crucially of all, entailing the truth of epistemological beliefs such as: *I know I am not dreaming, a brain in a vat etc.* It will be best, however, if I postpone further explanation of what sort of epistemology falls out of accepting a non-cognitive theory of perception until much later on (Chapter 6).

In contrast, a challenger of infallibilism is a certain kind of *fallibilist*: one who maintains that our best epistemic case for external world beliefs can *never be* one where their truth is guaranteed beyond all possible doubt.⁹ The hallmark of fallibilism is the insistence that no matter what we do to verify the truth of an external world claim via perception and perceptually-based processes, it will always be possible that we are still nonetheless deceived in some sense. Among fallibilists, one should distinguish what I will call *optimists* from *pessimists*. An optimist maintains that, despite our predicament, some of our external world beliefs can still be *epistemically justified*--justified, that is, enough so that we can call these beliefs knowledge as long as they also happen to be true. The matter of an external world belief's positive status as epistemically justified is not significantly enough affected by the fact that we are unable to *conclusively verify* its truth via perception and perceptually-regulated actions; so says the optimist. Among optimists, in turn, one should distinguish that of the *non-discursive* versus *discursive* variety. Non-discursive optimists maintain that external world beliefs are fallibilistically epistemically justified, in part, because the matter of being epistemically justified *does not* hinge on whether or not one can *argue effectively* against someone who maintains the following position: that our external world beliefs are not only false, but objectively improbable. Discursive optimists, on the other hand, hold their view because they believe that even though fallibilism is the proper position, we can nevertheless respond adequately to someone who maintains that our external world beliefs are objectively improbable, by actually performing the act of successfully arguing that some external world beliefs are (in some sense) likely to be true.

This all but throws a lot of the terminology that I will employ at the reader--perhaps too hastily. But of course the above distinctions will become clearer through the course of their deployment in context in the discussion that follows. Those familiar with contemporary epistemology will, however, probably have already noticed the similarity that lies between being a discursive versus non-discursive optimist and being a certain kind of (non-skeptical) epistemological internalist versus externalist. It is, in part, because the last qualification of "a certain kind" is indeed necessary to include when pointing out this similarity, that I have chosen to avoid this more familiar terminology. It is simply unclear which version of internalism, for example, is the best one to choose as the paradigmatic instance of the view.¹⁰ Indeed, the central reason why I have in many cases chosen sometimes cumbersome neologisms rather than familiar terminology in what follows, is that I would like to show, perhaps at the expense of greater readability for professional philosophers, that a version of the skeptical problem can be motivated that floats free as much as possible from a lot of related contemporary *debates* in epistemology. I would like to show that, at the very least, the *core issue* of skepticism is a relatively clear-cut, uncontroversially describable affair. If I am right one can, that is, motivate (and eventually address) the problem of skepticism in an *theoretically neutral* fashion--a plus for anyone who, like myself, is uncertain about how to even go about deciding who *in fact* holds the most plausible position in many of the currently raging debates in contemporary epistemology.

In any case, even for those who have a stake in those debates, the task of the first three chapters will have relevance. That task is to show that *no matter what kind of fallibilist* you are, a kind of skeptical-reasoning-induced paradox can be shown to arise

from your position.¹¹ In other words, first of all, if you are a fallibilist, then you ought to be what I will call a *pessimistic fallibilist*. This view maintains that the fact that our external world beliefs can't be perceptually conclusively verified implies, by way of a clear-cut line of reasoning, that to maintain that they are (likely to be) true is, to introduce another term of art, *presumptuous at best*. This results in a problem because we, as a matter of contingent fact, also believe ourselves to be in a strong epistemic position with respect to some of our external world beliefs. However, (in a qualified way that I will need to develop through the course of the first three chapters) this last belief is inconsistent with what fallibilism shows about the presumptuousness of our external world beliefs. According to this understanding of matters, *the problem of philosophical skepticism* is that there is something paradoxical that results from accepting fallibilism--namely, the tension between what we take our epistemic position to be with respect to certain external world claims, and what fallibilism (again, in a sense I will have to qualify) entails regarding that epistemic position.

The aim of the present chapter is to first present something that resembles the usual case for fallibilism and then partly motivate pessimistic fallibilism. I should stress the word "partly," however, because in this chapter I do not intend to argue against *non-discursive* optimism. The particular way that this sort of optimist responds to fallibilism, and how it differs fundamentally from both the discursive optimist and the pessimist's approach, is a matter that will be left unaddressed until Chapter 3. I will, however, try to show in this chapter why I think discursive optimism is misplaced.

In doing so, however, as I've already indicated, I will also try my best to spare the reader as much as possible from just another dry run in Humean-style epistemology.

Instead, the examination will be shaped by the following guiding consideration I will deploy the pessimistic fallibilist's reasoning in a way that is also as free from as many extraneous theoretical commitments as seems possible. In so proceeding, I strive for the kind of neutrality with respect to controversial matters that is, in my view, the defining characteristic of the skeptic's epistemological psychology. It is to that extent, I hope, a more than fair reconstruction that is also novel in several important matters of detail.

1.2 The Dream Argument¹²

Suppose that at the moment you were only dreaming that you are in the room you currently occupy, while in fact you were sleeping somewhere else entirely. Would you, in that case, be able to tell via any neurologically-based process that the statement *I am in my office* (or wherever you would want to say you were based upon the content of your dream) is false? Clearly not. The case is designed, after all, so that two important factors obtain simultaneously. On the one hand, things are going on inside your skin that result in your being *willing to assert that the external world stands a certain way rather than another*. But, on the other hand, things outside your skin are not that way. Clearly enough, whenever this combination of events occurs, any human subject will be unable to determine that his/her beliefs about the external world are false. Interestingly, moreover, what holds for a subject's external world beliefs in a case like that--a so-called *skeptical scenario* involving systematic deception--also holds for a belief like *I know I'm in my office* (or wherever). A subject in the above kind of predicament would be undergoing neurological processes that would make him willing to assert this (if, for example, we also asked him what he thought he knew about his current location); but would also be dead wrong.

Now for the all-important question that the reader, before proceeding, should ponder for a bit: *Can I be absolutely certain that I'm not that kind of deceived subject?*

Are you absolutely certain, that is, that:

(1) I know that I'm in my office

(2) I'm in my office

(or wherever else, as the reader's case may be) are true?

It might be thought that in order to answer this question with respect to (1), it is required that one *first* have a proper grasp of the concept of knowledge. This seems correct as far as it goes, though, I think, in a sense different from how a contemporary epistemologist might understand what it means to "properly grasp the concept of knowledge." In any case, since the goal at the moment is only to motivate fallibilism, we need not address what relationship answering (or being poised to answer) a substantive meta-epistemological question has with being able to properly answer the above question with respect to (1). In order to motivate fallibilism, it is adequate to try to show that for specific reasons that I will soon get to, one cannot be absolutely sure that certain *obvious necessary conditions* for knowing that one is in one's office are met. In other words, the fallibilist thinks that *any* worthwhile understanding of knowledge must answer the above question in the negative for both (1) and (2). Showing this requires only a rather superficial, and largely uncontroversial understanding of "our concept of knowledge;" sophisticated philosophical activities cannot, he thinks, prevent the result.¹³ Let us see how this might be done.

If you have been influenced by the usual philosophical treatment of this matter, then perhaps your thought proceeded in roughly the following manner when you

pondered the above question. In order for you to be absolutely certain that (1) and (2) hold, you would need to rely on what your sense experiences tell you. It is your sense experiences, after all, which make you willing to assert both claims, among other things, rather than something else (like that you are in a park and know you're in a park, rather than your office). Thus, for you to be absolutely certain of them, it seems that you would *already* have to have established that a certain relationship exists between the deliverances of your sense experiences and what is really the case. You would have to already know that those neurologically-based deliverances gave you the veridical result; i.e., that they were not deceiving you. But since the skeptical scenarios involving systematic deception are intelligible, it appears to follow that any human's reliance on neurological deliverances alone is at least *consistent* with continuously deceptively false information being fed to them while they rely on these deliverances.

And with this admission, whether you realize it or not, you have performed the philosophical equivalent of stepping on a landmine. On the one hand, given the kind of consistency just mentioned, you have certain things you do need to be able to *determine* are false in order for you to be absolutely certain of (1) and (2); you have to determine, for example, that it is false to think that your own neurological processes are continuously misleading you. But, on the other hand, all you have to go on to sort that matter out is more of your own neurologically-based deliverances. All you have to go on, then, is more processes that are of the very sort whose relationship with what is really the case at that time is what you need to determine. Thus, it seems right to conclude that you can never be *absolutely certain* that you are not continuously deceived in some systematic way, and hence that (1) and (2) are true. If you were absolutely certain that

both claims held, then you would at least be able to settle the matter, once and for all, of what the relationship between your neurological-deliverances and reality really was. You would, more specifically, be able to determine that your neurological deliverances cannot *possibly* be deceiving you in telling you that (1) and (2) hold. But, for the above reasons, you are unable to determine the latter. Call this the Dream Argument (DA) for fallibilism. With it, we can sum up the fallibilist position as follows. The fallibilist claims that we are not absolutely certain of (1) *because* we can come to determine, by way of the reflective procedure just outlined, that our sensory evidence for claims like (2) is inconclusive. To state the same view in a more condensed form, the fallibilist position is that by doing philosophy one can come to know that one doesn't know (2), and thereby deduce that one doesn't know (1).

1.3 Fallibilism Developed

I will subject DA to much scrutiny in chapters 4 and 5, which is where I give the central arguments for my own position. For now, I would like to take its conclusion as given and explore what it seems to force us to say about our overall epistemic predicament with respect to external world beliefs, as long as we don't opt for non-discursive optimism. (Thus, for the remainder of this chapter, by the term "sensory evidence" I will always mean that which is compatible with the truth of the skeptical scenarios).

With DA's conclusion, at least in my view, a paradox of a sort is born. The paradox is that, on the one hand, it seems somewhat ridiculous to claim that I am not *absolutely certain* that I am in front of a computer at the moment, *and* that I know that I'm in front of a computer at the moment. This is because I *see* the monitor that is a foot

or two before me, and I *feel* the keyboard under my fingertips, etc.; furthermore, I remember waking up, getting my coffee, and coming into the room to write. But, on the other hand, the above reasoning which motivates fallibilism also initially sounds convincing. It may even seem undeniable that *if* this were all a dream, then not only wouldn't I know the difference between that predicament and the normal veridical case, I could never discover what that difference might be given the perceptual tools I have available.

The question I would like to now address is this: What *concepts* (however roughly understood) does the fallibilist need, at a minimum, to make his epistemological point; and what does the fact that these concepts in particular seem necessary entail for his views on other theoretical matters? Performing this task will make just what the *pessimistic* fallibilist is maintaining become much clearer.

It is obvious that the fallibilist needs the concept of what I will call a *problematic epistemic possibility*. This is a certain kind of proposition; namely, one that a sophisticated enough subject may come to worry over when asking themselves about the strength of their evidence E for some belief that p. A problematic epistemic possibility is worried over when and only when a subject *first* comes to realize both that they fail to be absolutely certain that it is false, and that it is logically incompatible with the non-accidental truth of p; *and then* concludes on the basis of this dual realization that their evidence for a certain belief is inconclusive.¹⁴ (What I mean here by "non-accidental truth" will be explained, at least in part, very shortly.) I will henceforth refer to this as a problematic epistemic possibility's ability, after being worried over, *to reveal an epistemic flaw* in a certain subject's belief. A worried over problematic epistemic

possibility *proves* to a certain subject that his/her evidence for a certain belief does not, at that moment, guarantee that belief's truth. Notice, of course, that the bare existence of problematic epistemic possibilities is consistent with being able to subsequently come to establish a certain claim's truth conclusively by, for example, gathering more evidence. A fallibilist, however, holds the stronger view that all external world beliefs are *inescapably* epistemically flawed. They claim that all external world beliefs, even in ideal epistemic situations where we have gathered as much evidence as is possible, can be *proven* by way of DA to be supported by evidence that is inescapably consistent with problematic epistemic possibilities like, for example, that one is dreaming.

We can get more precise about what it is about these kinds of statements that makes them able to reveal epistemic flaws, by fleshing out what it means to be incompatible with the "non-accidental truth" of a certain belief. Certain problematic epistemic possibilities can do their job of revealing flaws even if they are also consistent with the truth of the flawed belief. Suppose I form a belief that a computer is in front of me on the basis of my sensory evidence. Suppose, e.g., that by "a computer" I mean here "any old computing machine." Next, suppose that the problematic epistemic possibility "Sophisticated scientists are stimulating my brain so that it appears as if a computer is sitting in front of me in an otherwise normal situation, while in fact my brain floats bodiless in a vat of nutrient fluid that sits in their lab" is pressed into service in order to reveal an epistemic flaw in my belief that there is a computer in front of me. This last statement's truth is *compatible* with a computer actually being in front of me. A computer could, as it just so happens, be sitting in front of my brain while it floats in the vat in the scientists' lab. This kind of problematic epistemic possibility, when considered alongside

my sensory evidence, therefore does not *guarantee the falsity* of the belief that I am in front of a computer. Nevertheless, such a statement still does its job of revealing an epistemic flaw because it is incompatible with the "non-accidental" truth of the flawed belief. Any facts about the external world which, in the relevant kind of deceptive case, would make my belief true would be facts that failed to play a *proper regulative role* in determining what I (the deceived subject) had going on neurologically which consequently made me want to say at that time that a computer was in front of me. Instead, what would be playing that kind of regulative role would be the actions of the scientists as they stimulated my central nervous system in the relevant way. What would make the belief that I am in front of a computer true in the above case (the fact that the scientists arranged the room in which my brain sat in a certain way) would be something (presumably) *sufficiently causally independent* from what actually regulated my neurological processes at that time, and hence what external world beliefs I possessed at that time. The relative spatial arrangement of the various macroscopic objects in the scientists' lab is a matter that might well vary while how my brain was stimulated was held constant (in the relevant sense), and vice versa. What this shows is that, according to the fallibilist, the precisely stated way in which my external world beliefs are epistemically flawed in general is that I cannot show or prove that what is regulating my current experience, and hence my external world beliefs about various environing objects, is not something *other* than simply the external world populated with those same mundane spatio-temporal objects.¹⁵

All of this implies that the fallibilist is committed to the view that the paradigm of an epistemically *flawless* belief is one where one's evidence guarantees its truth. What I

would like to do now is spend some time explaining what, minimally, a fallibilist must claim holds for this kind of flawless relation between one's evidence and a certain belief. By doing so I will shed light on what it is precisely that a fallibilist claims all external world beliefs *lack* evidentially--a matter crucial for understanding his position. Consistent with the guiding theme of this chapter, the goal will be to proceed in a way that frees the fallibilist from as many controversial theoretical commitments as seems allowable.

For *heuristic* reasons, though, I will now speak of one's sensory evidence as if it were only *propositional* in nature. Doing so is all but practically necessary, I think, for effectively introducing the following ideas. This should not be taken to suggest, however, that in my view a fallibilist is committed to their being nothing called "non-propositional evidence." As a matter of fact, in the next chapter I intend to show how a pessimistic fallibilist can side-step this issue altogether. But for now it is easiest to speak, if only *sotto voce*, as if one's evidence is *always* something propositional. While we are at it, as another heuristic that piggybacks on the first, let us also use the term *evidence* to refer only to those propositions that a subject has (i) *consciously considered* and (ii) *is absolutely certain are true*. Just like the claim that evidence is propositional, this should be treated at this point as nothing more than a simplification designed to merely help get certain ideas out effectively.

So constrained, one thing we can say is that an epistemically flawless belief must have an *objective import* that flows in particular from the fact that the evidence for it makes its truth *unassailable*. The central feature of being unassailable in this way is that of being a *truth-preserving inference*. The truth of the inferred belief "flows" directly

from the truth of the propositional evidence. Given its basis, in other words, the inferred belief just is immune from attacks that purport to reveal its falsehood.

To understand epistemically flawless beliefs, then, we must understand what is involved in this kind of *inferential unassailability*. Here I am primarily concerned with the question of *perceptual-experience-based* knowledge. Thus, it is perfectly allowable to henceforth focus the examination of inferential unassailability only on cases of inference where the inferred belief is a putatively *contingent* truth; something knowable, if at all, only *a posteriori*.

Given this qualification, plus the other guiding constraints of the current exposition, there is in fact a rather straightforward test for determining when a belief is epistemically flawless. It is a test that we can rightfully say the fallibilist believes is a powerful method for arriving at substantive epistemological conclusions. It is a test whose employment I see to be central to the persuasiveness of DA; specifically, to the all-important claim that it is *consistent* to think that we can have our particular neurological deliverances while also being continuously misled in some fashion. It involves the following reflective procedure. The subject, first of all, must be able to imagine or at least find intelligible what sort of situation the evidential statement(s) purport(s) to describe. They must also be able to do the same thing, more or less independently, for the inferred belief. After that, the subject must simply ask themselves if he/she can imagine or make intelligible sense of any situation that involved the conjunction of the evidential statement and the *denial* of the inferred belief. The fallibilist claims that an inference is unassailable for a subject if and only if as a matter of

psychological fact he/she cannot imagine or make intelligible sense of a situation describing such a conjunction, while meeting the above conditions as well.

The above procedure for testing for inferential unassailability results in a view of epistemically flawless beliefs that is appropriate given my desire for theoretical neutrality. It fails to engage many otherwise troublesome fronts, while also honoring certain basic requirements for strong evidential relationships. On the one hand, as I will explain below, it identifies as flawless those relations where the inferred belief follows *as a result of the laws of logic* from the evidence, as well as those where the belief follows (putatively) analytically. But, on the other hand, I think there is also a way to show that one need *not* be committed to the existence of logical truths or analytic truths (each supposed to be somehow distinct in kind from contingent truths) in order to be committed to the existence of epistemically flawless beliefs. *The fallibilist does not need to be committed to their being necessary truths of any kind whatsoever in order for him to make his epistemological position forcefully clear.* I will try to show this by first showing that there can be epistemically flawless beliefs that are *putatively* the result of an inference based on neither logical, nor analytic necessity. With this putative result in hand, I will then strengthen the case for a theoretically unburdened concept of inferential unassailability by both warding off certain objections as well as explaining certain theoretical virtues that are gained by pressing it into service.

Prima facie, e.g., if my evidence is *merely* "I'm in my office" and the inferred belief is "I'm not at home", then at least speaking for myself *qua* reflective subject performing the relevant sort of test, this describes an unassailable inference. Assuming that I *know* that I'm in my office, *from this knowledge alone*, I would claim, I am allowed

to unassailably infer that I am not at home. To see why, let us first completely disambiguate things and eternalize each sentence here so that they are indexed in helpful ways. So, the belief that I am in my office can be paraphrased as the belief that *K. H. is in room 254 of the English-Philosophy Building in Iowa City, IA as that building stood in October of 2009*. I claim, then, that knowledge of this alone (were I to have it at this time) would allow me to infer that *K. H. is not within the boundaries of the spatial location of 123 Somewhere St. in Wellesley, MA as that address was plotted in October of 2009*. Note that all the relevant criteria are met for this case. The inferred belief that I'm not at home--to hear revert back to the more condensed phrasing--is certainly of an, at best, *a posteriori* knowable sort. Secondly, I can make (in a manner that is determinate enough for the present purposes) intelligible sense of the situations that, as I understand such statements, would hold were I to properly accept both of them. Finally, I cannot make intelligible sense of *any* situation that would involve the correct affirmation of the evidence (so understood) and the correct denial of the inferred belief. There is just no way for me to even make sense of a situation where I am both at EPB as it stood in October of 2009 and also in Wellesley, MA.

Assuming (for the moment) the above is right, and the inference is unassailable as it stands, then notice that the relation between the truth of these two statements is of a weak modality when compared to that of strict logical entailment and analytic entailment. It is certainly not true that by *merely* understanding the meaning of the English phrase "I am in my office" one can discover that it is false that I'm at home. There are plenty of possible worlds compatible with the proper use of the relevant English phrases, where I am both in my office and also at home. So, the inference isn't obviously backed by an

analytic truth. Nor, most obviously, is there any sort of strict logical entailment involved in the inference described in such terms alone. If that inference is sound, in other words, it is not because of something that holds independently of an interpretation of the non-logical components of the inference. But the case passes the test for all that. Putatively, then, we have a case of inferential unassailability that is of the required weak modality.

Obviously, though, when assuming a higher-order perspective we can see well enough that there are a host of other beliefs of mine with which such an inference's unassailability must be consistent. I must also believe things like that when I'm in my office I'm not at home, that I can't be in two different spatial locations at once, and that my office and home are in two different spatial locations, etc. But what is far from obvious, I think, is that I must find the above inference legitimate only *because* I also accept these statements. *Prima facie* at least, the above methodology autonomously delivers the result. For one opting for neutrality, then, there seems to be no glaring reason to think that one needs to rely on there being anything like logical or analytic necessity in order to legitimate the inference. Why not simply leave the matter at that?

I suspect that one reason is that most philosophers will find it difficult to shake the idea that the above inference is either simply invalid, or else *enthymematic*. It is going to be difficult to rid oneself of the idea that I could only infer *once and for all* that I'm not at home in such a case if, along with knowing that I'm at EPB I also knew, for example, that

- (1) Whenever I'm at EPB I'm not at home.

It is only, it will be thought, *because* I also accept (1) implicitly that I am able to say once and for all that I'm not at home, when I also know that I'm in my office. If all I know to

be true is that I'm at EPB, the truth of the claim that I'm not at home is quite doubtful. Knowing merely that I'm in EPB seems to tell me nothing *one way or the other* about where my home is. And since it is also an ultimately contingent matter that I haven't decided, e.g., to hole myself up in some remote corner of EPB; then merely knowing that I'm in EPB seems *consistent* with my being at home. So, I must at least also know that EPB and my home are distinct locations, in order to unassailably make the above inference.

The above objection founders, however, on the fact that the fallibilist is free to use his knowledge that the above inference passes the relevant test as a means for *justifying* the truths of the statements that the above-described objector thinks must be presupposed. There is no better way to show that knowledge of a certain proposition p's truth isn't antecedently required to legitimate an inference than by showing that it is the legitimacy of the inference alone which can provide knowledge of p. I suggest, that is, that *one way* one can suitably paraphrase what one commits oneself to by maintaining (1), is as follows: "Assuming that I know "I am in EPB", the statement "I am not at home" will, on pain of its denial resulting in a situation that is unintelligible, be true". Of course, such a paraphrase should not be considered a *necessary* truth, nor the categorical that it grounds. And, even more obviously, there would be no sense to the claim that one *must* paraphrase all categorical statements like it in this way. But none of this matters. Dealing with the objection requires illustrating that there is nothing other than the notion of being *truth-preserving for a subject* that is required for understanding inferential unassailability, and hence for motivating fallibilism. Thus, the very purpose of making sense of the above as *one way* to paraphrase and also justify one's commitment to (1) is to show (eventually)

that the fallibilist can get their negative epistemological point on the table without committing themselves to the existence of things like, e.g., the laws of syllogistic logic. I am not saying I find the law-like status of logical truths in any way debatable. I am only trying to make clear that the fallibilist need not concern himself one way or the other with their existence in order to forcefully motivate his position.

Thankfully, aside from the above point--which is really nothing more than an act of digging in my heels--there is also a deeper, quite compelling reason for denying that the above inference *must* be considered enthymematic: it is simply not clear that there is something substantive that determines whether an inference is considered immediate versus enthymematic. One can, for example, take the inference rules of modus ponens and the material equivalence of a conditional and its contrapositive as primitive, and define modus tollens as a disguised case of those two rules.¹⁶ According to this understanding of what is/is not immediate, all cases of modus tollens would be considered enthymematic. But, importantly, one could just as well do the same thing, except replace the role of modus tollens with that of modus ponens (and include a rule of double negation).¹⁷ In this case all cases of modus ponens would have to be considered enthymematic. So, is it right to think that modus ponens is *really* modus tollens and contraposition and double negation in disguise, or (roughly) vice versa? Or, is the third option proper, and is it right to think that both inference rules are immediate? (There is no fourth option: obviously they both could not be considered enthymematic in the above-described ways.) I see no glaringly obvious way of deciding this, and I doubt that it even makes sense to say that there is a single right answer here. This provides all the more

reason, then, to think that it is misplaced to maintain that the inference described above *must* be considered enthymematic in order to be considered sound.

Perhaps it will be thought that the fact that statements/rules like modus ponens, modus tollens, double negation, contraposition, etc. *are each easily and--independently from one another--recognizable tautologies*, means that it is wrong to consider *any* of them enthymematic for the above reasons (thus opting for the third option presented above). But what is a tautology? At most I think one can understand a tautology in either of three ways. The first is what I would call a *truth-functional sense*. According to this understanding of tautology-hood, a statement is a tautology just in case it is a truth functionally complex statement where all possible inputs compute to a value of true. But this, of course, cannot be all there is to being a tautology. Tautologies are supposed to capture or represent independently understood *logical truths*. Merely being a truth functional statement that computes all inputs to the value of true is not enough to do so. Suppose, for example, we create *ex nihilo* the following unary truth function which I will call *wish fulfillment*. This truth function results in a true statement regardless of whether the single input statement's truth value that it takes as its argument is true or false. Clearly enough, the statement *It is a wish fulfilled that p* is, at most, a Pickwickian instance of a tautology!

This brings us to the second way one can characterize tautologies. One can understand them as statements whose *necessary truth* is easily grasped. Tautologies are statements that are *easily understood to be true in all possible worlds* (or else are built in an appropriate way out of statements that are so understood). But, obviously, given where we are in the dialectic one cannot simply deploy the notion of necessary truth here--or it's

related cousin of being true in all possible worlds--and expect to get away with it. So, how can someone who is trying to counter my above arguments by pointing to the fact that modus ponens, modus tollens, etc., are "independently easily grasped as necessarily true" effectively proceed? To my mind, the only thing they could do is rely on something at least rather close to what I describe above as the test for inferential unassailability, in order to flesh out this notion of being "easily grasped as necessarily true." What other options are available for describing what it *really* means to "easily grasp" a statement's necessary truth? Frankly, I have no idea what else it would be to "easily grasp" the truth of modus ponens or modus tollens except by performing that sort of test. Therefore, this understanding of tautology-hood only seems to *confirm* the fact that a fallibilist can get away quite easily with describing things in their preferred theoretically neutral sense.

The third way that one can understand a tautology is, not surprisingly, in both of the previously mentioned senses at once. However, I separated them for the purposes of dividing and conquering. I do not see why the fact that a professed tautology is both a truth guaranteeing truth function and an "easily grasped to be in every case true" statement means that modus ponens, modus tollens, etc., are obviously not *really* enthymematic for the above reasons. It is clear that the choice of truth functions that philosophers employ in, for example, the standard propositional calculus, is guided by a desire to in some sense capture an *independently understood* notion of logical truth. This is why they would block wish fulfillment's status as a tautology. So what does the work for granting a statement its status as an *independently understood* logical truth? The fallibilist opting for theoretical neutrality has an answer, and what follows from that answer is that to be a case of "easily grasped" modus ponens is not substantively different

from being a case of "easily grasping" that when I'm in my office I'm not at home. Merely insisting that only the former is "self-evident" does not do much to indicate a difference: it is clearly just another instance of digging in one's heels. Importantly, however, the one who wants to remain *neutral* here both has and, for the above reasons, I think, can easily maintain the high ground against any such objection.

As yet another way of strengthening the case for the idea that the objector's view isn't forced upon one who proceeds as I have above, I would like to now show that holding the view contrary to the objector's allows for a straightforward solution to the Lewis Carroll-style regress.¹⁸ This, of course, strengthens the case for the fallibilist's neutral view by presenting a nice theoretical virtue that stems from adopting it. It also shows how the view floats free of the problems that arise for one who wants to insist that a certain inference is/is not "really" enthymematic.

The Lewis Carroll-style regress arises, at least in one form, as follows: Suppose we abbreviate "I'm at EPB" as "p" and "I'm at home" as "q." If so, then, in order for me to move from knowledge of $p \rightarrow \sim q$ and p to the inferentially based knowledge that $\sim q$, I must also, it seems, know *that* $((p \rightarrow \sim q) \text{ and } p) \rightarrow \sim q$. Otherwise, why would I think that knowing both that $p \rightarrow \sim q$ and that p allowed me to conclude that $\sim q$? The only thing that could provide me with a sound basis for thinking this was allowed, it seems, would be my *additional* knowledge that $((p \rightarrow \sim q) \text{ and } p) \rightarrow \sim q$. But, of course, in order for me to understand *that* knowledge of p, $p \rightarrow \sim q$ and $((p \rightarrow \sim q) \text{ and } p) \rightarrow \sim q$ allows me to conclude inferentially that $\sim q$, I would, for similar reasons, also have to know *additionally* that $(p \text{ and } (p \rightarrow \sim q) \text{ and } (((p \rightarrow \sim q) \text{ and } p) \rightarrow \sim q)) \rightarrow \sim q$ held. But in order to understand that all of this knowledge allowed me to conclude inferentially that $\sim q$, I

would have to know that an even more complex proposition held, and so on. The problem, then, is that it seems impossible to understand how to get from knowledge that p and $p \rightarrow \sim q$ to $\sim q$, without presupposing that I have knowledge of the truth of an infinitely large set of increasingly more and more complex propositions. But this is just to admit that I have no idea how to get from the mere knowledge of p and $p \rightarrow \sim q$ to knowledge that $\sim q$.

One way to stop this regress at the ground level, however, is by allowing that the above test for inferential unassailability be a certain subject's *legitimation* for holding that $(p \rightarrow \sim q)$ --where the latter, in such a case, most accurately expresses the preservation of a subject's commitment to truth, on pain of the relevant sort of unintelligibility. Generalizing, one can claim that conditional statements in the context of modus ponens express the commitment, not to truths per se, but more accurately to *truth-preserving relations*. This would allow one to move from p to $\sim q$ *directly*, without the requirement for an additional premise $(p \rightarrow \sim q)$, and thus not allow the regress to arise at all. Put differently, in my view one can perform the above test at will for any pair of well-enough understood propositions, and thereby discover *whether* a host of relations of the form $(p \rightarrow q)$ hold given what situations one would commit oneself to upon hypothetical acceptance of the statements involved. If so, then we are able to say that our inability to make intelligible sense of a counter-example is what *legitimizes* the inference to q given *only* p . Thus, the best way to describe the above inference to $\sim q$ is simply that *assuming* knowledge that p , $\sim q$ has to be true. This is far different from saying that given that knowledge that p *and* knowledge that $p \rightarrow \sim q$, $\sim q$ has to be true. The inclusion of the premise $p \rightarrow \sim q$ is superfluous. In fact, on such a view, its truth (according, at least, to

one understanding of a statement of its form) could even be established *as a result of* knowing that given knowledge that p , $\sim q$ has to be true; something determined by whether the inference passes the above test for inferential unassailability.

Of course, to the extent that the soundness of the inference from p to some q would depend on an interpretation of the non-logical components of the inference, it would, by usual classification, not be able to be considered a logical truth. But I have already pointed out that not all inferentially unassailable relations express logical truths. Nevertheless, *if* we interpret the logical particle " \rightarrow " as being expressive of a truth-preserving relation (for a certain subject, on pain of the relevant kind of unintelligibility) then we easily extricate ourselves from an otherwise troubling theoretical difficulty. One can look at this as a virtue of the interpretation of " \rightarrow " and, to that extent, as confirming both that there are *sound* inferences of the above simplistic sort, as well as that the soundness of them does not rest on the fact that they are merely enthymematic.

The same strategy of appealing to the putative autonomy of the above test for inferential unassailability (on behalf of a party that wishes to remain neutral on otherwise controversial theoretical matters) can be employed in response to the claim that the inference is, in fact, enthymematic in other ways. Rather than belaboring the point I will show this for just one more instance. I admit that it *would* be perfectly intelligible to suppose that I am both in EPB and at home, *were* it perfectly intelligible to suppose that I can be in two different spatial locations at once. It might be argued, then, that the above inference passes the inferential unassailability test only because I also accept the claim that I cannot be in two different spatial locations at once. The real inference would be, in a much more fleshed out form:

- (2) EPB and K.H.'s home are in two distinct spatial locations
- (3) Any human subject cannot be in two distinct spatial locations at once
- (4) K.H. (a human subject) is in EPB
- (5) K.H. is not at home

Now, quite obviously the line of inference from (2) - (5) presents one proper way to come to know (5). But this in no way implies that it is the only way, and its presence alone as a sound inference surely does not show that moving directly from (4) to (5) in the manner described by the above procedure isn't also another way to come to know (5). I admit, of course, that a meta-linguistic examination shows well enough that I would not, as a matter of fact, reason in the latter way if I didn't also believe that my office and home are in different spatial locations, that a person can't be in two distinct spatial locations, etc. But what has not been shown is that I believe the inference to be sound only *because* I also have such prior knowledge. Consider, again, the analogous situation that pertains to accepting modus ponens, modus tollens, etc. Presumably my objector would agree that one would not find reasoning via modus ponens sound if one didn't also believe in the soundness of modus tollens, contraposition, double negation, etc. Whatever cognitive apparatus allowed one to grasp modus ponens "immediately" would, it seems, also allow one to grasp these other statements as well. Yet, my objector presumably would not take this sort of dependency to indicate that modus ponens is enthymematic. So, it is not clear why the same sort of dependency would require that I consider an inference from (4) to (5) enthymematic either.

Nevertheless, this objection in particular does point to the way in which I am hanging things rather heavily on the somewhat mercurial and psychologically-relative notion of what is versus is not "unintelligible for a subject". Is it *really* the case that I find it *unintelligible* that I could be in two spatial locations at once? Perhaps a clever

enough dialectician could press the matter, throw me into confusion, and cause a lot of the above to be unraveled. This must be admitted. But, of course, any problems with the notion of being "unintelligible for a subject" is, in the end, only to the detriment of the *fallibilist's* ability to motivate her/his position. As far as theoretical commitments go, the claim that sense can be made of what is versus is not "unintelligible for a subject" is something that the fallibilist must saddle themselves with, at least if I am right in claiming that the notion of *being truth-preserving for a subject* is central to her/his notion of being epistemically flawless. And it would be hard, at least for me, to make sense of the fallibilist's reasoning by way of DA in a way that didn't involve his trying to show that our external world beliefs are epistemically flawed. Furthermore, it is, I think, rather obvious that it is, at the very least, *by way of* a test for inferential unassailability that the fallibilist thinks (s)he can forcefully make her/his point. (S)he thinks it is obvious that we can each, after proper reflection, as a matter of psychological fact make *intelligible sense* of cases where our sensory evidence is what it is while any belief we might have about the external world is false. If the above way of looking at things does not describe how the fallibilist comes to understand his own epistemological predicament, then it is not clear, at least to me, what is supposed to make his position so powerful. If the notion of inferential unassailability is not robust enough, for example, to have *objective import* in the truth-preserving sense described above, then the fallibilist's position might even turn out to be philosophically innocuous.

It is also noteworthy that the above view of inferential unassailability quite easily accommodates cases of strict logical entailment, at least where the inferred belief expresses something contingent. If it is known that all dogs are mammals and that Fido

is a dog, and if from this knowledge it is immediately deduced that Fido is a mammal, then we clearly have a flawlessly justified belief. That Fido is a mammal is something knowable a posteriori, and there is no relevant situation *qua* counter-example that we can find intelligible here. A similar story can be given for how this view of inferential strength handles beliefs that rest on any purported *analytic* truths.

Let me now, finally, guard against the naturally arising worry that the characterization of inferential unassailability offered is too weak. It would seem to be if, e.g., it admitted something like a case where my known evidence was "Fido has a heart" and the inferred belief was "Fido has a liver" as a passable instance of an inferentially unassailable relation. Though it may as a matter of contingent fact be true that everything with a heart has a liver (just as it is may as a matter of contingent fact be that my office is not where my home is) it would seem to be too strong to say that knowing something has a heart gives one *incontrovertible evidence* (on pain of unintelligibility) that it has a liver. But obviously this sort of worry is doomed from the start. Examples like this do fail to be counter-example free. I myself can find *intelligible* a case where some Fido has a heart but fails to have a liver. Perhaps Fido is a newly evolved creature with something besides a liver proper; perhaps he is a freak of nature, etc. Given how I (and presumably the reader) might understand the truth conditions for "Fido has a liver"; it is *not* that the claim "Fido has a heart" has to be true on pain of otherwise describing an unintelligible situation.

In sum, what DA is supposed to show is that given the propositions which describe my current sensory evidence, I can find intelligible a situation in which those statements hold true but where the related ones describing the external world are false. A

fallibilist is one who insists that no truth-preserving relation is involved in moving from our sensory evidence to a related claim about the external world. Any reflective subject can come to see that fallibilism is true since they can, at the very least by way of the above procedure, come to see that all beliefs about the external world are epistemically flawed.

1.4 Pessimistic Fallibilism Conditionally Defended

The obvious retort to the fallibilist position as I have framed it would involve questioning his claim that our sensory evidence is intelligibly consistent with the denial of certain external world claims. But at least for now I will, however grudgingly, keep my promise and continue to grant that fallibilism is true. Now that we know what a fallibilist, at a minimum, needs to say about our epistemic predicament with respect to external world beliefs in order to make his point, I will now proceed to explain how fallibilism, so understood, leads to the position that maintaining that any external world beliefs are (likely to be) true is presumptuous at best. Why this result is paradoxical is a matter that will be developed in the next chapter.

As Barry Stroud has aptly put it, given that there exists problematic epistemic possibilities like "you are merely dreaming that p" we *lose the whole world* as far as our claims to evidential certainty are concerned.¹⁹ Nevertheless, at first blush this epistemological result may even seem quite palatable. DA may show that we cannot, given the sensory evidence we have, be *absolutely certain* that things are (roughly) as they appear. But absolute certainty seems to be the sort of attribution worthy only of belief in things like mathematical truths, tautologies, etc. Empirical propositions about the external world are of an importantly different, sometimes rather shakily-based sort,

and thus, it might be thought, may not require being subject to the same sort of high standard in order to be considered epistemologically sound to some appreciable degree. If the reader has this feeling, then certainly he/she disagreed with me above when I said that admitting fallibilism, by itself, resulted in a paradox. On this view, it is perfectly *natural* that our beliefs concerning the external world would *always* fall short of being guaranteed to be true by our sensory evidence. What would be paradoxical, I imagine this sort of person thinking, is if it were determined that our external world beliefs aren't reasonably held *at all*. The mere fact that they fail to be evidentially certain does not seem to establish this. Thus, the pessimistic fallibilist has some work to do if she is to show that fallibilism leads to some kind of problem.

The argument I will offer on behalf of the pessimistic fallibilist will proceed by way of first noting, somewhat in passing, the following view of hers (and, we will assume for the moment, the discursive optimist) which concerns what sorts of perceptually-based beliefs *are* flawlessly justified. Unlike external world beliefs, it seems that beliefs about our *subjective states* that we form as a result of our sensory experiences at a certain time are flawlessly justified by our sensory evidence. After all, even assuming fallibilism is true, merely by being one who undergoes certain sensory experiences I *can*, it seems, oftentimes conclusively verify that beliefs *about my own subjective sensory experiences* are true. Right now I am having a sensory experience that makes me willing to assert that there is a computer in front of me. It is, it seems, the simple fact that this sort of sensory experience, rather than another, is occurring which allows me to conclusively verify the truth of the following belief: *I am having sensory experiences as if a computer is in front of me*. If I reflectively ask myself *why* I think there is no possible way that I can be

misled into thinking that I am having sensory experiences as if a computer is in front of me, then the answer is simply because of the current sensory experiences I am undergoing at that moment. Importantly, even if, generally speaking, the worst were true and I *was* the victim of systematic deception with respect to external world beliefs, any experiential beliefs I might have would still be true. Unlike external world beliefs, when I consider the denial of the relevant experiential beliefs alongside my sensory evidence for them, the result *is* a situation that I do find unintelligible. The truth of experiential beliefs is thus something of which I can be evidentially certain (Further discussion of this admittedly thorny matter will be delayed until the next chapter).

Operating with a fallibilist picture, then, due to their failure to be able to be perceptually conclusively verified, external world beliefs have an epistemic status that must take a back seat to that of the content-related experiential beliefs. The fallibilist who wants to avoid radical skepticism in a *discursive* way--which, recall, is a strategy that attempts to *argue effectively* for the claim that even though fallibilism is correct, our external world beliefs are still in some sense likely to be true--thus has two options. Let me now mention and explore the first option at some length, and hold off even mentioning the other option until much later on. The first option I have in mind is to claim that external world beliefs, though not perceptually verifiable, are (in some sense) *known to be probable* based upon the fact that we can (in certain cases at least) verify the related experiential beliefs. It is to hold, for example, that *I know* that the fact that at the moment I *am* evidentially certain that I am having a sensory experience as if a computer is in front of me *makes it highly probable* that there is a computer in front of me.²⁰ It should be obvious why *knowing* such facts is something that the discursive optimist

requires. It is such knowledge, after all, which would be what they could effectively base their arguments against some challenger upon. (Without such knowledge, contrarily, it appears clear enough that their argumentative position would be too weak to support discursive optimism.) The first way that the discursive optimist might try to avoid radical skepticism while accepting fallibilism, then, proceeds with the natural enough idea that if I am to have any justification at all for my external world beliefs, it would have to stem from the fact that such beliefs are *known to be made probable* by my evidentially certain experiential beliefs. The purpose of this section is to explain why this familiar sort of move cannot work.

Of course, it is completely coherent to claim that there are non-deductive but nonetheless eminently *reasonable* inferences of the required sort. Take a case of a fairly run lottery where there are 1,000,000 balls in a big urn, but for which I have only bought one ticket representing one ball as a winner. In such a case, *if I know* that all of the conditions just mentioned obtain, then clearly it will be reasonable, to say the least, for me to infer that I am going to lose the lottery. Nevertheless, the inferred belief that I am going to lose is not evidentially guaranteed to be true in this case. Extraordinary luck could, for all I can tell, end up on my side.

A relevant question to ask at this point is: *When we actually make reasonable non-deductive inferences, in what precise sense can we say that the resultant belief is probable?* Answering this question in some detail will help to clarify what is involved in the project of determining whether any inference from experiential to external world beliefs is reasonable.

What I will now argue, again in a sense that I take to be on the *pessimistic* fallibilist's neutrally-minded behalf, is that the relevant sense of probability must be plausibly understood as a strong kind of epistemic probability that constitutively interprets the "makes probable" relation in terms of a propensity for producing true beliefs. To understand what particular notion of epistemic probability is required, we must remember that since we are exploring the prospects for discursive optimism, what we are most concerned with at the moment is whether or not under any circumstances the evidential certainty we have regarding the status of our experiential states can be used to *argue* for the claim that certain external world beliefs are probable. Clearly enough, then, the sense that our external world beliefs must be probable is epistemic: such beliefs must be probable given our knowledge of our experiential states. Moreover, the notion of epistemic probability must be understood in the following strong way: a belief that *p* is made epistemically probable by one's knowledge that *E* when and only when a subject *knows that E makes p probable* (in which case it follows, of course, that *E* also does in fact make *p* probable). Again, the requirement that the subject know that *E* makes *p* probable is in place because this is required in order for a *discursive* optimist to support their position.

My suggestion, to state the matter baldly, is that a pessimistic fallibilist will happily endorse this understanding of epistemic probability, but also interpret the sense of "E makes *p* probable" as follows: *E makes p probable when and only when knowledge of E makes any inference to p have a high propensity of coming out correct*. What I would now like to argue for on the pessimistic fallibilist's behalf is that if the discursive

optimist cannot show that external world beliefs are probable *in this sense*, then he/she cannot effectively argue against the relevant challenger.

Let us make things more concrete by supposing that the following information is known by some subject:

- (3) I have purchased a single ticket representing a single ball as the winner in a lottery of 1,000,000 balls
- (4) The lottery was run fairly [the winning selection wasn't rigged by anyone]
- (5) The winning ball has been selected
- (6) Only someone who has purchased a ticket can win the lottery
- (7) All the usual laws of physics that have held so far also held for the case in question. [This avoids supernatural ways of becoming a winner.]

This is simply another, more finely qualified version of the first described lottery case. It seems natural, right away, to suggest that in such cases the *known* truth of (3) - (7) is in some sense *epistemologically relevant* to the truth of:

- (8) I have not won the lottery.

We cannot say, however, that such a case represents the limiting case of epistemological relevance. We cannot say, that is, that the known truth of (3) - (7) *guarantees* (8). But we can, I think, somewhat provisionally say this: according to a certain appropriately understood metric of similarity, an inference from the known truth of (3) - (7) to (8) very closely approximates a case where (8)'s truth would be guaranteed. There is, after all, a situation where the inference to (8) would be guaranteed to be true that is very similar (in a sense to be qualified below) to the situation described by the inference from the known truth of (3) - (7) to (8). That situation would be where knowledge of (4) - (7) still held, but where knowledge of (3) was replaced with that of:

- (3*) I have purchased *no tickets whatsoever* representing any balls as the winner in a 1,000,000 ball lottery.

An inference from the known truth of (3*) - (7) to (8) would indeed be deductive. Since this is so, and since (3) represents a situation extremely similar to that represented by (3*), we can provisionally say that an inference from the known truth of (3) - (7) to (8) closely approximates a case of inference where (8)'s truth would be guaranteed. We can thus say that the truth of (3) - (7) is extremely epistemologically relevant to the truth of (8). Loosely speaking, (8) is almost guaranteed by knowledge of (3) - (7). To deny this, on this view, is self-contradictory for the same reasons that make it sound self-contradictory to deny that knowledge of (3*) - (7) completely guarantees the truth of (8). We can, in other words, discover by a reflective procedure that it would be absurd to deny that (8) is almost guaranteed by (3) - (7); a procedure that is of the same sort that allows us to discover the absurdity of denying that knowledge of (3*) - (7) guarantees (8)'s truth.

Making better sense of this sort of epistemological relevance thus requires making proper sense of the manner in which (3) and (3*) are similar. Generally speaking, similarity seems best understood as a *triadic relation* characterized by: X is similar to Y with respect to Z. Claims to the existence of a dyadic similarity relation are easily falsifiable. Being a whale is similar to being a dog in terms of reproductive processes (both give birth to live young), but not similar at all to being a dog in terms of primary habitat. So, saying that a whale is similar to a dog, if treated as a simple dyadic relation, is both true and false. The only way to remove the contradiction is to paraphrase the claim "A whale is similar to a dog and not similar to a dog" in a way that makes clear the triadic relations involved. Thus, we might say instead "A whale is similar to a dog with

respect to certain reproductive processes but not similar to a dog with respect to primary habitat."

This is helpful, because it relieves one of the burden of performing the hopeless task of explaining how (3) is dyadically similar to (3*). And since similarity seems to be a triadic relation, the natural idea is that it is only assessed relative to some sort of metric, i.e., an *ultimately stipulative* (though, as we will see, not arbitrarily) means of either measuring or comparing relative amounts of similarity between the relata. What metric one chooses for couching one's similarity judgments is going to hinge, first and foremost, on what sorts of things one is comparing and what one's goals are for the examination involved. If, for example, one is comparing things according to their size, then a natural metric to choose would involve comparing the objects' length in terms of some chosen unit of measurement, etc.

In our case, that of comparing (3) to (3*), the metric that we should choose also quite naturally falls out; especially given our goal of uncovering an epistemologically relevant relationship between the two kinds of inference involved. It should be one which will result in the arrangement of *relevant situations involving what a subject knows*, according to the level of evidential assurance that is thereby provided, on the information given, for (8)'s truth. The relevant situations, so compared, will be all of those which, in part, involve a subject making an inference, based upon their knowledge of (4) - (7), to the claim that they failed to win the lottery. What individuates any such situation from the others, in a way that allows for their comparison according to the metric we will employ, is clearly best understood as depending on the particular number of tickets that the subject knows they have bought for the lottery. On this sort of set-up,

what should fall out as what I will call the *upward limiting case*, according to our metric, will be when (3*) - (7) is known by our subject, making (8)'s truth outright deducible.

Obviously, then, the *downward limiting case* would, in contrast, be:

(3**) I have purchased a ticket which represents *every ball* in a 1,000,000 lottery as a winner.

(3**) - (7) guarantees the *falsehood* of (8), and thus is best placed at the complete opposite end of the spectrum from (3*) - (7), as far as its comparison, similarity-wise, to other situations of the sort dealt with here is concerned. Without going into further detail, it seems clear that such a metric will place (3) and (3*) as close as can be similarity-wise (here ignoring reflexive similarity, of course). Thus, according to a metric which is relevant to establishing *the level of evidentially-based assurance* that a subject apprised also of (4) - (7) can place in (8), we can see that (3) is indeed very similar to (3*).

Since the cases involved here are those that ask what I can consequently infer based upon what I *know*, clearly we have uncovered the sense in which knowing (3) - (7) makes epistemically probable the truth of (8). Anyone wishing to show how the evidential certainty of experiential beliefs makes probable external world claims should thus be interested in determining whether they can make sense of how the latter beliefs are probable in this sense.

What I will show below, however, is that we can just as easily "read off" from such examples the fact that beliefs of this sort have a *high propensity* to turn out true in cases that describe inferences of that sort. Before doing so, however, I should spend some initial time explaining what I mean by a propensity interpretation of probability.

This sort of view is often explained by way of contrast with a so-called *frequentist* view. On a frequentist interpretation, the probability of a statement's truth will hinge on

the objective facts of the matter surrounding the frequency of correlation for the events or properties cited by the statement. So, e.g., the probability of the claim "The next raven I see will be black" hinges, on this view, primarily on the facts of the matter regarding the number of ravens in my surrounding area that are, as matters objectively stand, black. The truth of a claim is not probable, on this view, if the properties involved (being a raven and being black, e.g.) are as a matter of objective fact only very weakly correlated. On this sort of crude frequentist view, it is the way the world merely contingently happens to be which determines the likelihood of a statement's truth.

Prima facie, this view is far from counter-intuitive. However--in a manner related to the so-called "Single Case Problem" that is often cited as a count against it--it is when one assesses how this sort of view handles the probability of statements that concern *the future* that one begins to see its overall inadequacy.²¹

Presumably, propositions that are assessed for likelihood can express rather specific truth conditions. They can refer, then, to matters that do not *directly* relate to broadly characterizable correlations that objectively obtain at a certain time. So, take the proposition that "This coin will turn up heads the next time it is tossed." This statement refers to a single coin in one's possession, and thus not to coins, even fair coins, generally speaking. Suppose furthermore that one has never, as it just so happens, flipped the coin in question. What could the frequentist's understanding of the probability of the above claim be in this case? They have no actual track record, no objective fact of the matter to appeal to that would allow them to determine the correlation between tosses of *that* coin and results that land heads. A crude frequentist seems forced to say that the probability that *that* coin will land heads when tossed is undefined. *This* should sound counter-

intuitive. There seems to be some perhaps loose way we can get a handle on the likelihood of the coin's landing heads in the trial yet to be undergone. For example, it seems intelligible to claim that it is more likely that this coin lands heads than it is that I'll win the lottery tomorrow. The intelligibility of this claim suggests, at least, that the probability of this claim about the future is something we can meaningfully discuss. If so, and if an interpretation of probability like the crude frequentist's implies that it would be undefined, then this serves as a good basis for a *reductio ad absurdum* of the view.

Perhaps, though, the frequentist can effectively refine their view as follows: they can point out that this coin is, we may reasonably assume, similar enough to most other coins in the sense relevant for assessing its likelihood of landing heads (in terms of its basic atomic structure, the uniformity of its density, etc.). They can then appeal to coin tossing cases in general, and thereby be able to claim that the future coin toss's likelihood of coming up heads *is* well-enough defined. This makes the matter of the future coin flip's turning up heads, on this sort of frequentist interpretation, hinge primarily on the frequentist understanding of the probability that "All flips of a fair coin will turn up heads" and then an instantiation to this coin in particular. The probability of the statement with the rather specific truth conditions involving my coin that is yet to be flipped thus hinges on the probability of the generalization of which the relevant trial here can, plausibly enough no doubt, be seen as an instantiation.

But, if so, then on a frequentist interpretation the probability of the general statement will hinge on the *facts of the matter* regarding all coin flips (involving coins of that sort) past, present and future. Such a frequency will obviously be yet to be born out in nature, referencing as *it* does trials far flung into the future. It seems, then, that the

frequentist is back to square one. On their view, to understand the probability (in the sense they have in mind) of a future coin toss one must appeal to future coin tosses of all fair coins generally speaking. This replaces one thing dimly understandable on frequentist grounds with another.

To deal with this, the frequentist is forced to say that my future coin flip's likelihood of turning up heads is a function of the frequency of all (similarly enough structured) coin tossing's landing heads *whatever that frequency may in fact turn out to be*. Unfortunately, however, this makes the frequentist's position no longer recognizably frequentist. A key (and familiar) epistemological point, after all, is this. At the time of the assessment of the particular coin flip with which we are concerned above, what the frequency of coin flips resulting in heads will be relative to the (possibly infinite) set of overall trials will be left underdetermined. Taking all of the *so far observed* coin flips as one's data, the relevant track record is consistent with an infinite number of different permutations of results that each conceivably could, as matters are born out, represent what is the case for the total frequency of head tosses for all coin tosses whatsoever (this assumes, furthermore, that such a frequency will even be well-defined in the instance where the size of the class of overall trials is infinite).²² Therefore, by appealing to a frequency that depends on how matters *will* play out, the frequentist is required to rest the assessment of the probability of my coin flip on the truth of a certain *counterfactual*. Namely, one which accurately describes how all coin flip trials *would* work out as far as the proportion of heads results is concerned, *were* all such trials to be performed. Given the above-described underdetermination, this is the only way a frequentist can speak about there being *a fact of the matter* regarding the relative frequency of results that are

heads for all coin flips whatsoever past, present, and future; and as we saw a frequentist needs to make sense of the probability of the next coin flip's landing heads in a way that depends on such a fact of the matter. A certain counterfactual *must* be true, in other words, in order for the frequentist's understanding of probability for such instances to be well-defined. But, of course, to rest matters of probability ultimately on the truth of a counterfactual is to no longer rest matters on an *actual* frequency. It results in a position that is, as I see it, no longer recognizably frequentist. In sum, the problem is that frequentists need to be able to make sense of the probability of events yet to happen, but since the only correlation-dependent way that sort of probability can be understood is one that requires that a certain counterfactual be true, it follows that a frequentist interpretation of the matter can't be correct.

When considering how to make sense of the probability of a future event, then, we are pushed in the above way towards resting that probability on the truth of a certain counter-factual. This is just what a *propensity* view, at least in what I take to be its clearest form, does.²³ A propensity theorist still is like the frequentist in claiming that probability is an objective, mind-independent property of the world. It is just that what determines the probability, rather than any actual frequency of results, is instead the physical make-up of the relevant system which gives rise to a particular *tendency towards a certain result's obtaining*. Thus, on a propensity view, as it just so happens any fair coin flip (be it past, present, or future) has the likelihood of producing heads that is, say, .5. But this is simply because, as matters shakeout, coin flips of that sort have a *tendency* to produce heads 50% of the time. One way to express the point as a propensity theorist, though not the only way, is that a propensity measures the *tendency* that a certain set of

initial conditions has for producing a certain limiting relative frequency of a certain outcome, as the number of relevant trials approaches infinity. An important point, however, according to the propensity theorist, is that this tendency need not be born out precisely in what is the fact of the matter regarding such a limiting relative frequency. The tendency that a flipping of a fair coin under normal gravitational conditions, etc., may have is to result in a limiting relative frequency of heads results, as the number of trials goes to infinity, which is $1/2$. But this does *not* mean that as matters play out in nature this limiting relative frequency will approach exactly (.5) as the number of trials moves towards infinity. Perhaps the actual limiting relative frequency of heads results approaches (.49), or something like that. On the propensity theorist's view, however, this would not change the fact that the results tend towards (.5). Contrast this with the above, refined frequentist view, which would identify the likelihood of a certain trial with the fact of the matter regarding such a limiting relative frequency. So, in the last-mentioned instance, a frequentist would identify the likelihood that the next flip of a fair coin was heads with (.49), assuming this was the fact of the matter. The problem with this last view, recall, was that in order for the frequentist to think there *is* such a fact of the matter, it seems they must be committed to the truth of a certain counterfactual like: *Were the number of trials to be extended to infinity, the relative frequency of heads results to overall tosses would approach x/y .* We can see how this sort of commitment naturally suggests a propensity view. The propensity theorist should agree, at least in my view, that such a counterfactual's truth is required. They just would add that what this true counterfactual indicates is an experimental set-up's tendency toward a certain limiting relative frequency, as the number of trials moves towards infinity. So, the likelihood of

my yet to be flipped coin's turning up heads is related to the tendency that a certain experimental set-up has for producing a limiting frequency of heads results as the number of trials approaches infinity. I say "related to" here because propensity theorists do disagree over how to understand propensities: for example, whether we should *identify* the probability that a certain event has of obtaining with the relevant propensity.²⁴ Nevertheless, all propensity theorists agree that it is a certain tendency or disposition that certain events have for obtaining, given particular sorts of situations as initial conditions, that is relevant for assessing probability.

Of course, as far as philosophical clarity is concerned, explaining propensities in terms of the truth of counterfactuals of this sort simply explicates one poorly understood notion in terms of another.²⁵ I have nothing to offer here to help clarify this matter on the propensity theorist's behalf. It does seem to me, in any case, that the intuitiveness of the problems with a frequentist view reveals the implicit acceptance, on our part, of just such a counterfactual (as long as one is like me in thinking that it makes perfect sense to assign probabilities to future trials). We think it makes sense to talk of a future coin flip's likelihood to occur, I would argue, because we implicitly accept that *some* such counterfactual is true.

The relevance that this brief foray into interpretations of probability has to the current discussion is this: I would argue that to the extent that implicit commitment to a propensity view explains why we feel it makes sense to say that the probability of *future* events is well enough defined, it *thereby* explains why it should make equal sense to say that the probability of a *hypothetical* event can, in certain cases, also be well enough defined. Since our gut feelings support a propensity view in the coin flip case, then I

would argue that they should do so for the latter case as well. In both cases, after all, we are not resting our assessment on how things actually are. We are instead assessing probability in terms of what result a certain "experimental set-up" will have the tendency to produce. It is just that, for hypothetical cases, it is what one might call a "*thought* experimental set-up" that is at work, rather than what might very well be an actual experimental set-up such as the staging of a fair coin flip.

With this in mind, the lottery examples I have examined above do seem to support the following counterfactual: *Were one to examine the result of a large enough set of situations where it was known that (3) - (7) held, then one would find that an inference to (8) came out true (to put it crudely) much more often than not.* A situation of the sort where (3) - (7) are known and (8) is consequently inferred, in other words, *tends towards the result* that is a true belief in (8). We can say, in other words, that it is the hypothetical "experimental" set up of the knowledge of the conditions (3) - (7) which determines the tendency that the psychological event of inferring (8) has for producing a true belief. The considerations that guide our assessment here, I would argue, are exactly parallel to what guided how to assess the probability of the future coin flip. However, since the case involving the reasonable inference from (3) - (7) to (8) appeals, for its "experimental set-up" to things known by a subject, we are likewise right in affirming that a case like this involves something that is epistemically probable. In my view, speaking on behalf of a fallibilist opting for theoretical neutrality, the best way to understand what it means for one's knowledge to make the truth of a certain inferred claim likely is that inferences of the relevant sort tend towards the result that is a true inferred belief. Another way to put the point is that, one cannot, it seems, reference the propensity that inferences from the

knowledge of (3) - (7) to (8) have for producing true beliefs without also referencing the directly related (and perhaps even identical) epistemic probability of the inferred belief, and vice versa.

Now, with this characterization of the sort of probability at work in a reasonable non-deductive inference having been given, let me now explain what is relevant as far as is concerned the attempt to show that external world beliefs are probable given their evidentially certain experiential bases. If the above analysis is on the right track, then knowing that external world beliefs are epistemically justified requires knowing *what it is about situations involving knowledge that an experiential belief is true which is highly epistemologically relevant to the truth of some inferred, content-related belief about the external world*. If it can be shown that we *lack* this last sort of knowledge, then this will show that one natural avenue for the discursive optimist fails.

So, e.g., suppose I form a belief about the location of my favorite book upon the basis of testimony of someone whom I have never met. Suppose I know, however, that the place I'm in has an equal distribution of completely honest people, habitual liars, and those somewhere in between in terms of truth-telling habits (please indulge me here). My relevant evidence in such a case is thus only this bit of (under-described) background information, and the known testimony of the person involved. Let us then suppose that I inferred that my book is where they say it is. In such a case I obviously have no clue one way or the other what trusting a claim made by that person has to do with producing a true belief. For all I know, a situation of this sort is identifiable with (and thus similar to) a case where I was told something by an habitual liar, all but guaranteeing the falsehood of the produced belief (this, with some further qualifications, would be the

relevant downward limiting case). For that matter, the inferential situation could, for all I know, be identifiable with one where the person is an honest and very observant individual, all but guaranteeing the belief's truth (here, again with some further qualifications, we would have the relevant upward limiting case). Or, of course, the case could be identified with one that fell somewhere between these two extremes. The point is that I here have no way of determining where to place the situation on the relevant similarity metric. This means that what I know is compatible with being identified with *any* kind of case found anywhere on the relevant spectrum (i.e., the upward limiting case, downward limiting case, and anywhere in between). For this reason, and this reason alone, I suggest, it would be utterly wrong to think that one could ever argue for the claim that this kind of belief is likely to be true. As I will often subsequently put it, this is because considerations of the sort just offered allow me to realize that believing in the relevant claim's (likely) truth is *presumptuous at best*.

So, moving now to the matter of the non-deductive reasonableness of external world claims, the only way I could know that the truth of "I'm having sensory experiences that p" made likely the truth of "p," would be if I somehow knew what it was about situations in which I performed the relevant inference which made them similar, in the relevant sense, to a case where an inference to p was guaranteed to be true. But, much as in the above case of trusting the above subject's testimony, there is simply no way to determine this given merely what we are evidentially certain of (at least if the fallibilist is right). The culprit for being unable to determine this is none other than the epistemological lesson of DA. According to DA, for all we know the situation in which it seems to me that p is not only similar but outright *identifiable* with the downward

limiting case where my evidence is consistently misleading. Perhaps, for that matter, it is only similar to such a case, and thus identical with one where, say, I am prone to be provided with misleading evidence 90% of the time. Perhaps, lastly, a certain instance of sense-experience that *p* is identifiable with the relevant type of case where my belief that *p* is *guaranteed* to be true. Unfortunately, I have no clue how to sort this matter out. The point, of course, is that once the truth of my sensory evidence is shown to be cut off from the truth of any external world claim (by way of the relevant reflective procedure) it will follow that it will be consistent with a seemingly vastly disparate number of "experimental set-ups," so that the actual propensity that my inferred belief concerning the external world has for coming out true is anyone's guess. Much like the previous case, it seems clear that to say that my external world beliefs are (likely to be) true given merely what I know regarding the content of my experiential states is *presumptuous at best*, and therefore unreasonable according to a discursive optimist's way of looking at matters.

I don't think it is too hasty to claim that all falliblist attempts to defend the inferentially-based non-deductive reasonableness of external world claims will fail for the above reasons. The lack of a deductive-style connection between our evidence and the truth of such claims forces us to admit the consistency that that evidence has with all sorts of inferential situations where the inferred belief involved is even *guaranteed* to be false. It is fascinating to remember that what got this whole epistemologically pessimistic ball rolling was the seemingly benign admission that the skeptical scenarios are intelligible.

1.5 Wright's Argument Against Radical Skepticism

So far we have shown how being a fallibilist along discursive optimist lines fails, as long as one tries to base the probability that an external world claim is true on the fact that a related experiential belief is known with evidential certainty. There are, of course, other fallibilist-friendly ways one might attempt to ward off radical skepticism. One obvious way, which I will address in chapter 3, is to adopt a form of non-discursive optimism. But, not surprisingly, there are also plenty of other discursive optimist avenues one can take. I would like to end the chapter by critically examining one interesting proposal given by Crispin Wright.

Wright, in effect, tries to show that radical skepticism is a *self-defeating doctrine*. He proposes the following line of argument. If one accepts that it is *unreasonable* to hold any beliefs concerning the external world that are formed on the basis of sensory evidence alone *because* of the fact that one (it is granted by Wright) knows nothing of the *cause* of those experiences, then it follows that it will be unreasonable for the same general sort of etiologically-related reasons to hold any beliefs that are based upon the *operations of one's intellect*. But if so, then since the sort of position we are worried about here is clearly one arrived at via argument--an operation of the intellect to be sure--it will follow that that position itself is unreasonable. So, Wright tries to argue in his (1991) paper that any argument against discursive optimism is self-defeating. If it is sound, then we can show that this would entail that we should not trust the sort of reasoning which leads to the very conclusion it purports to reach.

The argument proceeds as follows:

[W]hether or not dreaming as ordinarily conceived suspends them, sound intellection--understanding, inference and reflection--is, like perceiving, subject to

[etiological] constraints...following a proof is a specific form of understanding, and the claim to have done so is consequently answerable to what the subject is subsequently able to do in the way that any claim to understand is...There is, for example, no absurdity in the idea of a subject who, while capable of grasping each of the ingredient thoughts involved in ratifying a sophisticated proof, lacks the ability to follow the reasoning involved; yet can nevertheless rehearse it, with every confidence and a strong sense of familiarity, as a result of hypnotic suggestion.” (p. 105).

The point here is, in a certain sense anyway, well-taken. There is a difference between, say, parroting a certain deductive inference, and actually understanding that it holds. Part of understanding that it holds would clearly require understanding why the truth of the premise is relevant to the truth of the conclusion. In other words, knowing, say, that the premises of a Barbara syllogism entail its conclusion requires understanding how, on the information given, there is no imaginable way for the premises to be true and for the conclusion to be false. Merely being able to parrot a certain instance of such an inference does not require having this sort of understanding. Nor for that matter would, for example, being able to parrot the inference as well as parrot an explanation, while also failing to understand how that explanation is relevant. Loosely speaking, we can say that *part* of what is involved in *understanding what one utters* is that one have a workable enough idea of what sort of situations (and/or, depending on the utterance, other claims) one is (or, if one is speaking hypothetically, would be) committing oneself to obtaining by *correctly* holding the relevant claim. That one does or does not have this kind of workable idea is something that will become manifest, at the very least, in one's subsequent behavior. One who is merely parroting a certain deductively valid inference, and even parroting a related explanation, for example, will not know what to say when subsequently questioned about, say, the relationship between the explanation offered and the original inference. That someone knows more or less what they are talking about can

most times be manifested in the fullness of time through merely continuing to speak with them. Wright's point expressed above is well-taken, then, in the sense that we must recognize that there are, for example, certain *methods* for coming to assert that a particular entailment relation holds that are better suited to understanding it than others.

With this claim granted, Wright then points out that radical external world skepticism rests on the claim that

there are no conclusive indications by which waking life can be distinguished from sleep...[which in turn rests on the] compelling thought that experience cannot disclose its own causal provenance as part of its proper content...But the idea is no less compelling for episodes of thought. That, as a train of thought develops, the ingredients are caused in certain ways cannot itself be *manifested* by their collective content (though it may, of course, be part of it). Descartes, when he ventured to regard cogent intellection as marked off by phenomenological characteristics of clarity and distinctness--whatever exactly they are--missed an insight whose counterpart in the case of perceiving he seemingly did not miss. For the fact is that episodes of apparently cogent intellection, no less than episodes of apparent perceptual experience, may, for all that is phenomenologically evident to the subject, have an [etiology] inconsistent with their being genuinely intellectual/perceptual (p. 106).

The reader can probably see where this is going. Anyone motivated by radical skepticism in a manner that depends on pointing out our inability to know *what causes* our perceptual experiences, will, Wright argues, have to in turn admit that the same sort of skeptical considerations can apply to merely reflectively-based beliefs. If so, and if the inadequacy of discursive optimism is a position that is supposedly reached by reflective considerations, then the reasonableness of this same position will be called into question by such considerations, at least if we accept the truth of the premises designed to motivate it. So, Wright goes on to claim, we cannot accept the truth of the premises designed to motivate any view that denies discursive optimism, since they are self-defeating in this manner.

I take it that "to understand" is what Ryle would call an "achievement verb" (Ryle 1949, 149). If so then one does not understand a proposition to varying degrees, but either does or does not understand it. The way to respond to Wright is thus to first explain how understanding a line of reasoning is construable as understanding a (sometimes rather complex) statement. On this way of looking at things, to be able to cogently follow a certain line of reasoning is nothing more than being able to understand a sometimes rather complex statement. So, e.g., one could construe understanding the validity of a Barbara syllogism as understanding that the following complex statement holds:

- (8) Its being true that all A are B and that f is an A guarantees the truth of the claim that f is a B.

It is by no means a stretch to take simply understanding (8) as a paradigm case of cogent intellection. Now, as I understand Wright, he is pointing out that there are etiological constraints on being able to understand (8). If one affirms (8) merely as a result of hypnotic suggestion, for example, then Wright's suggestion is that one has not really understood it. But that is the important point; being a case of cogent intellection is, for this reason, simply identifiable with being a case where one has understood a statement such as (8). So, though we may impose etiological constraints on understanding (8), there is, it seems, no way that one can understand it and yet fail to have undergone cogent intellection.

Wright's only option, then, is to say that if we accept the fallibilist line of thinking that denies discursive optimism, then this somehow commits us to skepticism regarding whether or not we can ever successfully understand something thought to hold as a result of a reflective procedure. But what I will now argue is that to say this requires that Wright stretch the meaning of "to understand" in ways that are, well, hard to understand.

Take what is perhaps the easiest case to prove: that of understanding when a negative logical relation holds. For example:

(9) Being a case of seeming to see that p does not entail that p.

Now, a perfectly conclusive way of indicating that one understands that (9) holds, it seems, would involve offering a series of relevant counter-examples. One who understands (9), for example, would be able, when prompted, to present a slew of cases of systematic deception as counter-examples to the claim that its seeming that p entails that p. If Wright is to deny that being able to do this conclusively indicates one's understanding that a logical relation such as (9) holds, then he seems to have a rather peculiar sense of "to understand" in mind. It is utterly bizarre to continue to question whether someone has "really" understood (9) after they have been able to cite innumerable counter-examples as their explanation for why they think it holds. Wright's overly schematic point about the necessity for a proper etiology does not stand up to the obvious fact that how someone explains their reasons for maintaining a claim can, in cases like the above, be a *conclusive* testing procedure for determining whether or not they have "really" understood it.

And, for the purposes of showing why radical skepticism of the brand presented in this chapter is unmotivated, the irresistible retort here of questioning the idea that this last-mentioned point is an "obvious fact" is a clear losing strategy. The way that radical skepticism was motivated above involved a *first-person* enterprise. It involved someone coming to understand that the plausibility of fallibilism alone makes maintaining that our external world beliefs are (likely to be) true presumptuous at best. As a first-person enterprise, therefore, it will render irrelevant the sort of facts that admittedly might (to

state the extreme position) forever prevent one person from conclusively determining whether *another person* has "really" understood something. It can be granted for the sake of argument, in other words, that in many cases it will be even impossible for a conclusive determination of this last sort to take place--perhaps because we cannot "peer inside the head" of the other person. So, one can grant for the sake of argument that it is false that the above point is an "obvious fact" in any *intersubjective* type of case. But now try to apply the same type of concern to the first-person case. What would it even be for one to *merely think* one has conclusively verified that one has understood a claim like (9) *because* one has made sense of counter-examples, while in fact being wrong on this very point? No concerns can be brought to bear here by *another* person. You are the only person involved, and you are the one who has made sense of the relevant counter-examples. In sum, it is plausible to say the least to maintain that making sense of counter-examples in the first-person case simply *constitutes* understanding a claim like (9). Wright's generically stated worry over etiology does not stand a chance here.

Another set of instances of cogent intellection that are somewhat easy to indicate when one has versus has not "really" understood them, are those which involve *positive affirmations* of logical relations that depend only on affirming syntactically construable relationships between propositions according to the meaning of the logical particles alone; thus (8) above.

How, in such cases, can we conclusively indicate that we have understood that a statement or proposition of this sort holds? Well, one way, manifested indeed in our subsequent behavior, is to proceed to actually produce an isomorphic physical representation of how the relevant information hangs together (perhaps, for example, by

drawing an appropriately constructed Venn Diagram). After doing so, we can literally *see* that the information expressed by the conclusion is already expressed in the relevant way by the premises. By noticing this rather simple and obvious isomorphism between how the information of the premises and conclusion is in fact related and how our representation depicts matters, we can indeed understand that to deny the truth of (8) would require that we deny that the picture is itself, which simply doesn't make sense. In other words, if one and only one type of Venn Diagram, for example, pictures both the information provided by the combined premises as well as, simultaneously, the information provided by the conclusion, then it seems that the only way to deny that the truth of such premises guarantees the truth of the conclusion (and thus deny (8)) would be to deny that that picture is itself! We can conclusively understand, it is hoped (!!), that to deny the latter is absurd, hence we can understand that we cannot deny that the relevant logical relationship holds. On this view, being able to conclusively understand that certain positive logical relations hold hinges on being able to conclusively understand that a particular instance of the law of self-identity holds. Here, the latter law applies to the picture which is isomorphic with the logical relationship that holds between the information provided by the combined premises and that which is provided by the conclusion. Now, to speak frankly, it just seems rather hard to follow the suggestion that anyone might affirm the law of self-identity only *mistakenly*. If Wright is to rest his argument against radical external world skepticism on his own brand of skepticism over whether we really know the latter law to hold, then I think this shows his position to be rather unconvincing. It is allowable, of course, to attack these issues in the extremely generic way that Wright does. What seems clear, however, is that when we get down and

dirty and discover what particular sorts of statements we need to be committed to in order to say we conclusively understand logical truths, we find that these statements are rather hard to deny. Those who think we can trust in our conviction that the laws of logic are cogent can take heart in the fact that it seems that the best Wright can do to respond is to appeal to abstracted remarks about how, *in general*, etiological constraints apply also to cases of cogent intellection. When what follows from this remark is placed alongside the law of self-identity, I think it is clear that the latter wins hands down as far as plausibility is concerned. If Wright is to deny that we can be really certain that we understand that the above instance of the law of self-identity holds, then it is at best not in any way clear what he has in mind.

This leaves us to consider one last case of cogent intellection, which is admittedly the trickiest to defend. I have in mind the various sorts of affirmations of the truth of so-called analytic propositions, as well as the other kinds of unassailable inferences I have made much of above. Such propositions are different from those like (8) because analyzing them syntactically alone does not reveal their obviousness. In such cases the syntactical structure would be something like, say, " $p \rightarrow q$ " or " $(x)(Bx \rightarrow \sim Mx)$."

So, e.g., the usual thought is that the old standby:

(10) All bachelors are unmarried

is an example of such an affirmation. The present question is whether sense can be made of cogent affirmations of statements like it. The matter is indeed particularly pressing since some cogent affirmations of this kind are inextricably involved (albeit somewhat implicitly) in coming around to the pessimistic fallibilist's position. For example, affirming something like the following is required:

- (11) All instances of non-deductively reasonable inferences are cases where the known truth of the premises makes probable the conclusion.

We must be at least implicitly willing to affirm something like (11) if we are to come around to understanding how fallibilism leads to philosophical trouble regarding the reasonability of our external world beliefs. If so, we had better get a handle on the difference between a cogent case and a non-cogent case of affirmations of this type.

Thankfully, to defend the idea that we can understand well enough what makes for a cogent version of such an affirmation, I need only appeal to my earlier discussion of unassailable inferences. The requirements for understanding cases of such inferences are flexible enough to include what are or seem to be rightly called cases of analytic entailment, and thus the requirements are flexible enough to understand precisely when cogent affirmations like this have taken place.

Though, indeed, it is correct that an examination of the syntactical structure of purportedly analytically true statements is not enough to reveal their unassailable truth, it is still no accident, I think, that such statements have a *conditional sign* as their main connective. One will *never* come across an analytic statement that has the syntactic structure of simply "p", for example, at least where further analysis of it will not unearth the presence of a conditional as the *main* connective (or of what can, through well understood stipulative definitions of logical particles like "v" and "|", etc., be translated as a case of such a centrally-placed conditional). Indeed, we can see an attempt to translate such statements into the language of quantification theory, were it to be undergone, as an attempt to achieve a natural syntactical understanding of such statements, *whereby* a "→" is appropriately centrally placed (or able to be so placed via stipulative definitions). Thus a natural translation of (10) comes out as something like $(x)(Bx \rightarrow \sim Ux)$, and (11) might

come out as something like $(x)(y)(xNy \rightarrow xPy)$. Given this last rendering it is natural to take the values of the variables to be at least meaningful natural language statements and even perhaps propositions. The relational term "--N--" stands for something in English like "The statement --- renders non-deductively reasonable via inference the statement --" (something different would have to be involved here if we were quantifying over propositions), and "--P--" something like "the truth of --- renders probable the truth of --." My point here is only that such a rendering in quantification theory would have such a statement turn out false only when what precedes the main connective " \rightarrow " is true and what follows it is false.

Of course, such a rendering alone does not allow us to say what we can for *logical* truths, namely that such a statement will come out true regardless of the interpretation of the non-logical components of the statement. There are plenty of consistent interpretations where the quantificational renderings of (10) and (11) given above come out false. To that extent, then, we obviously have not captured their unassailability merely via such a rendering in terms of quantification theory. But what we have done is understood them in terms that allow us to perform the test for inferential unassailability *on their English interpretations according to our current grasp of that natural language at a certain time*. We can thus understand a cogent version of such affirmations to be one where the joint affirmation of the antecedent and the denial of the consequent would result in a situation that is unintelligible given how we *antecedently* understand the relevant English phrases. Of course, Wright's brand of skepticism could then morph into the question of whether we have "really" understood the relevant English phrases. But this is to change his argument. Skepticism over whether we really understand the relevant

component English phrases would not be a case of questioning the cogency of the proper sort of intellection; or, at least, not straightforwardly. The kind of intellection Wright needs to be able to be conditionally called into question involves, rather, understanding complex statements like (8) - (11).

In sum, the fact that we can make clear enough sense of cases of cogent intellection such as understanding statements like (8) - (11) forcefully disarms Wright's argument. If Wright is to question the explanations given above for when we really can understand statement of this type, then it seems that at that he has a peculiar idea of what is involved in understanding them. And, importantly, as long as it is clear enough that we can understand statements of the relevant sort, pessimistic fallibilism can be motivated in the above fashion.²⁶

1.6 Conclusion

This concludes my examination of the logical connection between fallibilism and a certain kind of skeptical position. We have found, for one, that the fallibilist only needs a notion of inferential unassailability that holds the notion of *truth-preservation* as central, in order to show why external world beliefs are epistemically flawed. I was able to use the result of that examination to help explain why Wright's argument is inadequate, and I will make further use of it below for similar negative purposes. I also showed that the interpretation of probability that one can see as at work in the pessimistic fallibilist's arguments is a kind of robust epistemic probability that also makes use of a propensity theory interpretation of "is probable."

As the next order of business, I will now turn to a more in-depth examination of the most obvious philosophical reaction to accepting the arguments given above, that of

adopting and developing the epistemological approach of what I have called pessimistic fallibilism. This will pave the way for explaining what I think is wrong with non-discursive optimism--a matter that I will address in chapter 3. Aside from aiding in that explanation, defending certain implications of pessimistic fallibilism will also prove to be imperative for explaining and defending the positive position I offer in chapters 4 - 6. Most crucially of all, the specific form of infallibilism that I defend in those chapters will also rely on the view that being in a strong epistemic position with respect to certain claims hinges on being able to refute various kinds of challengers. It is thus necessary to make clear that the full implications of this general sort of position are quite within the realm of plausibility.

CHAPTER II PESSIMISTIC FALLIBILISM DEVELOPED

2.1 Introduction

Pessimistic fallibilism is best seen as the philosophical position that is without qualification concessive to the kind of position motivated in the previous chapter. That position is that fallibilism entails that maintaining that external world beliefs are (likely to be) true is presumptuous at best. The following analogy explains well enough what I mean here. That is, if fallibilism is true then the relationship between our sensory evidence and any external world claim is exactly like the relationship between the evidence base comprised of "I am taking part in *some sort* of lottery" and the belief "I'm going to lose the lottery." It should be clear enough that in the latter case one doesn't have even the slightest clue how likely it is that one is going to lose the lottery, given the paucity of one's evidence. What I have hopefully shown (following Hume) is that the same goes for any external world belief, as long as fallibilism is true. Further pursuing the philosophical program suggested by pessimistic fallibilism requires, then, among other things, determining what to make of the epistemological status of *any* of our sensory-experience-based beliefs, given that all of this seems to clearly follow.²⁷

However, we must also remember that the pessimistic fallibilist' approach aims at theoretical neutrality. As a result, in finding central *only* the fact that the above sort of argument is compelling, it is absolutely crucial to emphasize that the pessimistic fallibilist's approach to epistemology diverges radically from one guided instead, even in part, by the desire to properly understand epistemological *concepts* like justification and knowledge. This last kind of investigation seeks an analysis of (for example) being epistemically justified that is (ideally) counter-example free, where whether or not an

offered proposal has a worrisome counter-example depends on its logical relationship to our intuitive, from the gut grasp of the relevant concept. It is well-known, however, that providing these sorts of analyses is a project mired in insuperable difficulties.²⁸ The pessimistic fallibilist floats free of such difficulties. He can admit, then, that the concepts he employs are even extremely counter-intuitive in the relevant sense. He can admit to this because developing a position that jibes well with our intuitive grasp of justification, etc., is not something that he is interested in doing. Instead he is interested in developing a view that responds appropriately to the fact that the above arguments appear so convincing. It is a characteristically *philosophical* reaction to the apparent fact that once fallibilism is granted, the claim (roughly stated) that external world beliefs are presumptuous at best follows. How we in fact use the term “justified”, etc. across a *broad* spectrum of cases has no direct bearing, it seems, on the success of this sort of enterprise. This is not to say, however, that the pessimistic fallibilist denies that a more or less explicit and more or less worked out understanding of what look like central epistemological concepts is required in order to forcefully motivate his view. But any work the pessimistic fallibilist puts forth on such matters to motivate his position is meant to serve, at best, a context-dependent *demystifying* function. It is supposed to make his claim *that a kind of paradox arises that results from the relationship between what we continuously naturally take our epistemic position with respect to certain external world beliefs to be and what fallibilism entails regarding that epistemic position* clear enough to one who does not understand initially what precisely he has in mind. Obviously, one can perform an act of context-dependent demystification like this without also claiming that what one has done is provided an interesting, non-circular synonymous rendering of any

mysterious phrase found within the relevant claim. So, the pessimistic fallibilist would not be bothered to learn that what *looks well enough like* his talk of *being justified*, for example, can be shown to diverge radically from the extension suggested by our intuitive grasp of that concept. He does not think, for example, that, intuitively, being justified is synonymous with being able to refute various imaginable kinds of challengers. But he does nonetheless think that the claim that external world beliefs are presumptuous at best quite obviously follows from fallibilism, and that this is due only to our recognizable failure to be able to perform one such act of refutation. (Further development of the pessimistic fallibilist's position regarding the relationship between being unable to refute a certain kind of challenger and what follows about the status of *external world* beliefs in particular, will be postponed until the next chapter, where I criticize various non-discursive optimist views.)

As a result of so proceeding, however, what does warrant fundamental stressing throughout the philosophical development of the position, if it is not in any remote sense whether the implicitly employed understanding of justification (for example) accords with our intuitive grasp of that concept? If we cannot, for example, fault the pessimistic fallibilist for employing an unintuitive understanding of epistemological concepts, then how *can* we find philosophical fault with his position? Relatedly, what makes the position worrisome or problematic if it is not, even in part, the fact that it makes use of an intuitively plausible understanding of things like knowledge, having good evidence, etc.? To answer these questions we must bring more sharply into focus just how the process of motivating pessimistic fallibilism occurs. If we wish to formulate and develop an epistemological position that results only from the fact that a persuasive, theoretically

neutral argument for the position that external world claims are presumptuous at best exists, then how we proceed in formulating this general epistemological position will depend on how that argument proceeds. Only once this process of argumentation is thoroughly examined can it be made clear enough that forcefully motivating the pessimistic fallibilist's position *does not* depend in any way whatsoever on whether the epistemological concepts employed therein have a broadly intuitive application; or whether any epistemological concepts do for that matter.

2.2 Assertion and Negative Epistemological Arguments

The most salient general feature of the pessimistic fallibilist's argument is, perhaps, the fact that it is a *negative* epistemological argument. How, then, do negative arguments of this sort proceed? They do so by challenging *in a certain context* the claim that one has sufficient evidence for a particular assertion that one has made, or is willing to make. There are no negative arguments of the relevant sort offered unless there are assertions challenged. In order to understand the pessimistic fallibilist's general epistemological outlook, then, we must understand what it is about the very speech act of assertion that allows for negative epistemological arguments to gain a foothold. It is true that one way to describe what we are up to here is an examination of *intuitive* cases of assertion. But let it *not* be thought that the examination to follow depends upon quibbling over the central uses of the term "assertion." Another way of describing what we are up to runs as follows: an examination of what results when a certain kind of challenger is successful in showing that one is in a certain kind of weak epistemic position with respect to something that one has said, or is willing to say is true; thereby contradicting what one implicitly or explicitly thought was the case regarding one's epistemic position with

respect to the claim in question. If the reader thinks "assertion" intuitively picks out speech acts different from those of the sort in question, then I invite them to henceforth translate accordingly by replacing my use of "assertion" with their own favored term.

I think that the relevant feature of the speech act of assertion is this. When we make assertions we *contextually imply* that we find ourselves in a favorable epistemic position with respect to what we are asserting.²⁹ Assertions have what might be called their own *momentum*. Merely by saying something with a certain amount of force in certain context, one can dig oneself into an epistemological hole. This is made evident enough, I think, by noting the infelicity involved in the following utterance: "There *is* soda in the fridge *and* [notice that here I did not say "but"] I don't think there is any good reason for believing that there is soda in the fridge." The reason this utterance sounds odd, I would claim, is because its net cognitive *cum* pragmatic content is zero. What the first conjunct forcefully puts forth, the next takes back in full. This helps to confirm that making assertions takes a certain amount of both guts and preparedness. We do not make assertions, *especially* in philosophical contexts, without having something to back them up.

With this, we can present the following development of the pessimistic fallibilist's position, which addresses the paradox that results from the apparent success of his arguments. It is clear enough that *the only reason the pessimistic fallibilist's negative arguments should be considered interesting is because most people are, as a matter of fact, continuously willing to assert that things stand in the external, mind-independent world one way rather than another*. The pessimistic fallibilist's position that all external world beliefs are presumptuous at best would not be considered problematic, it seems, if

there weren't those of us around who are willing to make assertions like this. Notice, after all, that maintaining that a belief is presumptuous at best when the relevant claim is not, as a matter of fact, backed by our willingness to assert it results in a position that is either uninteresting or obvious. It is uninteresting when we are unwilling to make the relevant assertion *because* we lack beliefs about our favorable epistemic position with respect to the claim in question as well as its negation. I am not bothered by the fact, for example, that I cannot refute someone who maintains that at the moment the number of people in New York City is not even. To point this out is uninteresting, I think, because of the simple fact that I am not at the moment willing to assert that the number of people in New York City *is* even. I don't, as a matter of fact, at the moment take myself to be in a favorable epistemic position with respect to that claim. When, on the other hand, it is pointed out that we can't refute someone who claims something that we are *already* willing to assert, the position so fleshed out is not so much uninteresting as obvious to us (take, for example, a line of argument which showed that I cannot refute someone who denied that $2+2=5$). Perhaps even at this point, then, it is clear why developing the pessimistic fallibilist's epistemological position does not require the employment of concepts meant also to have a broad intuitive appeal that extends through logical space. It is simply a fact, made evident by even the most superficial examination, that people are (and probably always will be) willing to make assertions that depend on the external world being one way rather than another--that we all continuously take ourselves to be in some positive epistemic position with respect to such claims. In my view, it is *just* this contingent, everyday fact that the pessimistic fallibilist exploits in order to show that a

kind of paradox arises. In my view, moreover, *the problem of philosophical skepticism* is nothing more than the fact that this paradox does seem to arise.

To further develop the pessimistic fallibilist's view, we must note that he thus has a way of dividing *simple* assertions into three categories based upon their epistemological status; where by "simple" I mean here assertions that do not have inferences as their content (how "inferentially complex" assertions are dealt with by the pessimistic fallibilist is a matter that we have already at least partially addressed in the previous chapter). Among assertions like this that *survive* skeptical challenge, there are those of either the *epistemically flawless* or *non-epistemically flawless* variety. What determines which of these last two types of assertion/belief can serve to survive a skeptical challenge are, I think, the following (implicitly or explicitly) contextually-fixed factors. The matter hinges both on whether there are any probability judgments that the disputants agree are warranted, as well as on how good the disputants agree one's evidence must be in order to survive the relevant skeptical challenge. The only way that a non-epistemically flawless belief can survive a skeptical challenge is if i) the inquiring skeptic and the one willing to make the relevant assertion agree that certain epistemically prior probability judgments are warranted; and ii) the (implicitly or explicitly) agreed-upon standards for how good one's evidence must be to meet the challenge are low enough. When both of these conditions are met, it will follow that certain assertions about what grounds one's belief can be deemed adequate enough to settle the matter. For example, suppose you and I agree (as most people do) that under normal conditions visually-based beliefs about one's immediate surroundings are highly likely to deliver true beliefs. Suppose, further, that the contextual standards in place only require the belief in question to be highly likely to be

true in order to survive a skeptical challenge. Then, clearly, I could meet a challenge offered by you by pointing out that I am sober, my eyes have just been checked by a doctor, and that a certain belief was visually-based (and I am assuming here that whether or not a belief was visually-based is something easily decidable in the relevant context).

Now, obviously, showing instead that a belief was epistemically flawless would be more than enough to quiet such an obliging fellow as the above-described challenger. But if no such background probability judgments are agreed to be warranted, then it is clear that (aside from, perhaps, some controversial cases)³⁰ the only way to quiet a skeptic about the epistemic soundness of a simple assertion is to give *conclusive* reasons for thinking that the belief is true *simpliciter*. The only alternative would be to think that we could give non-conclusive reasons for believing something that were also of the sort that, unlike those mentioned above, didn't rely on already agreed upon probability judgment(s) for establishing the point. But it is not clear (again, aside from perhaps some controversial cases) how this could be done. The only sort of non-conclusive (high-probability-determining) procedures that exist seem to be ones that rely, always at their last step (albeit almost always implicitly), on the existence of an already agreed upon probability judgment as the deciding factor--such as that visually-based beliefs under normal conditions are likely to be true.

But, to repeat, in cases where we are bracketing such judgments, it is quite clear that any methods we employ for effectively countering various kinds of challengers must be ones which attempt to verify conclusively that the statement in question holds. And at this point we arrive at a crucial realization. Namely, the *only* way to conclusively verify a statement under such circumstances would be to make some use of (among other things)

the way things are at the present moment as a crucially deciding factor. In cases where the reliance on probability judgments of the relevant sort is considered illicit, it is much more crucially the way that an agent makes use of how things actually stand around him at the relevant time (or, as the case may be, within him--more on this below) which can, in certain cases, provide the grounding needed to survive a certain sort of challenger's challenge.

Of course, even when the matter hinges on whether an assertion can be conclusively verified via a certain procedure that makes use of how matters stand around one, there must be certain background judgments that the disputants agree hold. For example, they must agree that a certain method works as a decision procedure for settling the matter of the assertion's truth value. But this is agreement over the epistemological worth of a particular kind of procedure, and not agreement over the truth value of a certain probability judgment. Only for the former sort of agreement does the *eventual* result of the decision procedure take priority, so that (sometimes) who is right can be quite up in the air until those results are in. In some cases, furthermore, the disputants agree at least initially that the matter *can* be settled once and for all by the relevant investigation. For the moment I will hold off addressing the important question of whether, in some such cases, it is ever worthwhile for the original challenger to go on to question the following higher-order matter: whether, that is, a certain decision procedure is *in fact* decisive in the relevant way. But we should note that there clearly are certain contexts where even the higher-order claim that a decision procedure works can be verified outright; namely, contexts where a certain meta-decision procedure is adequate to establish the point.

An obvious skepticism-induced regress threatens here, but at this point one should only notice that matters are quite different for the case where what is agreed upon are instead certain probability judgments. In such cases, it should be obvious to the disputants that what allows for a certain claim to sometimes survive a skeptical challenge is something that can neither be verified conclusively nor even non-negligibly confirmed at the present moment, given the other available epistemic practices. At the time of arguing over a visually-based belief, for example, it should be clear that no one will be in a position to verify that the following holds: *All visually-based beliefs formed under normal conditions are likely to be true.* Matters at least seem rosier for a statement like: *That is a cup.*

Although I have so far presented matters in a third-person fashion, let it not be thought that I've made the mistake of developing the pessimistic fallibilist's position in a way that is biased towards an intersubjective, "common sense" way of looking at epistemological practices. Simply put, the assertion-maker/challenger relation is reflexive for those who are reflective. Via a reflective enterprise, we can and often do challenge our own assertions. Getting an epistemological investigation going in this case is even much easier, it seems, than it is in a multiple party case, because some of the conditions mentioned above are trivially met. When we are involved in a reflective enterprise we do not ever have to first *agree with ourselves* that certain probability judgments hold/do not hold; nor that the standards for effectively answering the one skeptical (us) are low enough. Talking this way smacks of nonsense. What goes on, I think, is that we just employ these sorts of judgments and standards more or less implicitly, or we fail to rely on any such claims at all and seek instead to conclusively verify the matter. The above

account is thus flexible enough to accommodate both third and first-person epistemological examinations.

Along with the sorts of epistemological beliefs recognized above, the pessimistic fallibilist maintains that there are what we can here call *epistemically deficient* beliefs, which are simply those that demonstrably fail to survive an offered challenge. And with this we can now employ a helpful simplification. Let us simply stipulate henceforth that the only *non-epistemically deficient* sensory-experience-based assertions worth discussing are those that one is adequately prepared to *demonstrate* are flawlessly justified. Doing so is helpful not only to better understand a pessimistic fallibilist's general epistemological approach. Seeing whether there can be any assertions like this will also be important for determining the viability and overall force of my own infallibilist arguments that are designed to resolve the problem of philosophical skepticism. This is because *demonstrating in a certain context* that external world beliefs are epistemically flawless is, I think, the only viable way to resolve that problem in a realist-friendly fashion. We should see, then, what exactly it means to perform this kind of demonstration. That is, let us now ask: *What would it mean for a sensory-experience-based belief about contingent reality to be demonstrated in a certain context to be epistemically flawless?*

Speaking at a general level, to show a certain kind of challenger that a certain belief about contingent reality is epistemically flawless requires addressing a higher-order epistemological matter. It requires having an *impeccably* worthy answer to the question: *What grounds do you have for making your assertion?* If one cannot answer this question in a way that *proves* that the belief is epistemically flawless, then as a result of

the draconian standards so employed, one will be forced to withdraw one's assertion. Equally clear, however, is that the only answers that suit in this case are answers that one is, in turn, equally willing to justify forcefully. If this were not what held, then the sort of infelicity talked about above, that was involved in saying "There is soda in the fridge and I have no good reason..." would--again, as a result of the draconian standards so employed--become manifest. This pertains to what I prefer to call the *momentum* of assertions. To utter something with the force of an assertion, for example, is *always* to commit oneself to *more* than just what one has audibly stated. It is to contextually commit oneself to asserting *something else*, namely something of a higher order nature that bears upon what one takes to be one's epistemological position regarding the initial claim in question.

In this way, however, the pessimistic fallibilist's position engenders a kind of *complexity regress*; one that is akin to a related regress often pinned to a meta-epistemological position sometimes called *strong access internalism*. I will now briefly explain how the regress works for the strong access internalist. According to this position, very roughly put anyway, *generally speaking* a subject S is justified in believing that p only if S is justified in believing that S is justified in believing that p. Obviously, however, because this last view is meant to apply generally, it would have to also place the same sort of necessary condition on S's being justified in believing that S is justified in believing that p. In other words, given the proposed generality of the view in question, it follows that S is justified in believing that S is justified in believing that p only if S is justified in believing that S is justified in believing that S is justified in believing that p. But the same sort of necessary condition applies to the latter even more complex belief;

and it is clear, moreover, that the relevant kind of iterations will continue to be generated *ad infinitum*. In sum, a consequence of this view is that S is justified in believing that p only if S is justified in believing an infinite number of increasingly more complex beliefs, where the relevant kind of complexity results from the continued iterations of the sort just introduced. But, the thought may be, it seems impossible for anyone to even *comprehend* the more complex iterations, let alone be justified in believing them. For this reason, strong access internalism seems to imply that no one is justified in believing anything whatsoever; which sounds absurd. For this reason, this view is often rejected as misguided. In response, it is thought that we must somehow weaken the requirements for being justified.³¹

In a manner at least superficially similar to the strong access internalist, the pessimistic fallibilist maintains that, generally speaking, when one makes or is willing to make an assertion of the sort we are concerned with here, one contextually implies that one finds oneself capable of making a higher-order *assertion* that bears upon what one takes to be the strength of one's epistemological position for making the initial assertion. But by contextually implying that one finds oneself capable of making that kind of higher-order *assertion*, given the generality of the pessimistic fallibilist's view regarding what is involved in being willing to assert things with the relevant kind of force, it follows that when making a simple assertion one also contextually implies, in turn, that one finds oneself capable of making an even higher-order assertion, etc. Thus, according to the pessimistic fallibilist's view, one should never make an assertion that is supposed to survive a certain kind of challenge regarding its epistemic flawlessness unless one is

equally willing to make an infinite number of ever-increasingly-more complex assertions that each are claimed to be able to survive that very same sort of challenge.

But what I will now argue is that this presents no issue in and of itself. I will show, in other words, that the pessimistic fallibilist's position does not lead to the radical brand of skepticism that strong access internalism *may* indeed engender. There can in certain cases be beliefs that are shown to be epistemically flawless in the relevant sense.

Right away, however, we should take note of the following qualification that, it is hoped, will make my goal here seem much more attainable. The way I will deal with this issue is *not* by providing a general miracle recipe for showing any claim whatsoever one takes to be epistemically flawless to be epistemically flawless, in a way that is able to deal with the above problem so as to effectively silence any imaginable challenger. One of the things that my desire for neutrality effectively brings out, I think, is that in order to understand the resultant problem of philosophical skepticism, we need *not* think that there are claims that are reasonable/dogmatic *simpliciter*. Instead, we are only required to recognize, first of all, that there is good sense to the idea that we are in a certain kind of strong epistemic position with respect to a claim *when* we can effectively argue for it against various kinds of challengers. The problem of philosophical skepticism arises, then, because of the mismatch between what fallibilism entails our epistemic position to be with respect to external world claims and what we take it to be (given that we are all willing to make various assertions about the external world). The paradox arises as a result of what takes place *in a particular context*, then, one where we philosophically reflect in the manner already outlined. No claim about being dogmatic in maintaining things about the external world *simpliciter* needs to be defended in order to forcefully

present the relevant problem. For the very same reason, then, in what follows no procedure for showing any claim to be epistemically flawless *simpliciter* will be defended; proceeding as such is equally unnecessary. Instead, I will show how one can demonstrate a claim to be epistemically flawless *in a context where various (no doubt helpful) things are assumed*. The *relevance* of the fact (assuming it is a fact) that a claim can be so shown to be epistemically flawless in that kind of context will then be discussed at the end of the chapter.

An illustrative example will help things along nicely by, among other things, fleshing out just what sort of higher-order assertions *can adequately serve in a certain context* as what one contextually commits oneself to when making a more basic assertion. This will only show, of course, that there *can be* certain contexts where the pessimistic fallibilist can deal with the complexity regress issue. To reiterate, this is not to say that there may be other kinds of contexts where the higher-order issue requires the pessimistic fallibilist to hold his tongue.

Suppose, then, that I am a competent English speaker with normal visual faculties and that while staring at a red patch of fabric directly in front of me under normal lighting conditions, I form the belief that *I am having a red sensation*. Let us assume, furthermore, that the sense-experiences I had at that time *prompted* this belief, and that this prompting is in some central sense the result of the fact that:

- (i) I have learned to speak English proficiently;
- (ii) I am paying attention to the relevant portion of my surroundings;
- (iii) generally speaking (under the same type of conditions) I am able to distinguish red patches of fabric from fabric that is sufficiently differently colored (although relevant, let us here ignore issues having to do with vagueness);³²
- (iv) what is going on at that time is that I am looking directly at a red patch of fabric.

Despite my above qualifications, it should be acknowledged here that the obtainment of any or all of the conditions (i)-(iv) is certainly something that, on a particular occasion, might be open to legitimate skeptical attack. However, for the moment we are *not* concerned with dealing with such challengers. So, let us assume (helpfully, to be sure) that the subject in question is absolutely evidentially certain that conditions (i)-(iv) obtain.

Importantly, the sentence "I'm having a red sensation" is an *occasion sentence*, which is just to say it is the sort of sentence that I would assent to only *while* I was having sensations that are members of the set of sensations that would normally prompt my assent to such a sentence in the first place. Sentences like this contrast with *standing sentences*, which may elicit my assent even after I am no longer undergoing the sorts of sensations that would prompt assent to them in the first place.³³ (So, a statement like "The *Times* has come" is a standing sentence, because I would in many cases assent to it even after having sensations other than, e.g., seeing it flop on my doorstep.) For any occasion sentence, there will be a clear enough distinction between that sentence and the *set* of sensations that *would* prompt one's assent to it. Referring to the case at hand i) clearly what is happening to me as far as my sense-experiences are concerned is something that I would call sufficient for assent to the sentence "I'm having a red sensation;" but what is also true is that ii) a host of other experiences could, I recognize, just as well serve to elicit assent to this same linguistic item (suppose, for example, I was staring at a ripe tomato instead of a red patch of fabric). Keeping this in mind, for the case at hand, *the willingness to assert outright* the following claim (1) seems to properly

capture a subject's implicit commitment to the right sort of *epistemologically relevant* relationship holding between the relevant linguistic and non-linguistic items:

(1) *The truth conditions for "I am having a red sensation" are met at the present moment.*

I think one who is willing to assert outright that they are having a red sensation can claim that they are contextually implying, in turn, that they are willing to assert outright at least something close enough to (1). But much like the strong access internalist's higher-order requirement, this instantiates a *generally applicable* feature of being willing to assert something outright. This means, of course, that in being implicitly willing to *assert* (1) outright it is, in turn, contextually implied that one is willing to assert the following claim:

(2) *The truth conditions for (1) are met at the present moment.*

But this means that by being willing to assert outright that one is having a red sensation one *also* contextually implies that one is willing to assert that a statement even more complex than (2) in the relevant respect holds, and so on *ad infinitum*.

This should give the reader a proper feel for the kind of complexity regress that is engendered. But I see no reason to claim that this complexity regress present a problem for the pessimistic fallibilist. Notice, first of all, that unlike the regress of the strong access internalist, the pessimistic fallibilist's regress avoids requiring justified higher-order beliefs *in the justification* of lower order beliefs. The assertion of (1) differs *starkly* from, for example, the belief that I'm justified in believing that I am justified in believing that p in that it lacks epistemologically-relevant terminology (terms like "know," "justified," " etc.). Instead, commitment to the claim that one is flawlessly justified in believing (1) is implicit in *the very act of asserting outright* that one is having a red sensation. Nor is this

lack of explicit mention of epistemological terms in any of the relevant assertions an artificial element of the pessimistic fallibilist's position. It points squarely, yet again, to what I've been calling the *momentum* of assertions. Merely by performing the speech act of assertion we commit ourselves to having a certain strong epistemological position vis-à-vis the claim in question *without ever being required to explicitly say that we are so committed*. In this way, the pessimistic fallibilist avoid the aspect of the strong access internalist' is viewed that seems to lead to trouble. The problematic aspect that is avoided is the inclusion of epistemological terms in the higher-order claims themselves.

We can now flesh out the way in which the pessimistic fallibilist avoids the problematic aspect in question as follows. For the pessimistic fallibilist, there is a kind of *transparency* involved in the logical relationship between what one commits oneself to at the higher versus lower order that is (conditionally anyway) quite lacking for the strong access internalist's position. Regarding the latter's claims, it is not immediately obvious what it would even be to be, for example, justified in believing that one is justified in believing that one is justified in believing that p. How does the structure of justification work for higher-order claims? Depending on how we answer this question, it may turn out that, for example, the conditions that must be met in order to be justified in believing that we are justified in believing that p are *far different* from the conditions that must be met in order to be justified in believing that p. (Perhaps, for example, being justified in believing that we are justified in believing that p requires *doing philosophy*, while being justified in believing that p merely requires use of what are *in fact* reliable faculties.) In any case, my point is just that a substantive question must *first* be answered here in one way rather than another, before the strong access internalist can sit well with her higher-

order requirement. In stark contrast, the pessimistic fallibilist requires higher order claims to be epistemically flawless that are *easily understood to be logically equivalent* to the lower order claims. It is just immediately obvious that, for example, for any occasion sentence it is true at a certain moment that it is true at a certain moment that is true at a certain moment that p when and only when p, and so on. No substantive epistemological question must first be answered, then, in order for the pessimistic fallibilist to make good on the claim that we can understand the higher-order assertions when and only when we can understand the most basic belief. If we can understand the claim that p in the case where p is an occasion sentence, then it is already obvious that we can understand the claim that it is true at a certain moment that it is true at a certain moment that is true at a certain moment that is true at a certain moment...p as well.

Ultimately, it is the fact that this kind of obvious logical equivalence exists between the higher order and lower order claims, *plus* facts having to do with what I have been calling the momentum of assertions, which explains why the pessimistic fallibilist's complexity regress presents no problem. Given the relevant kind of transparency, baldly stated, the pessimistic fallibilist has, I think, a simple way to address the issue. He can claim that the occurrence of the actual process of sensory experience that prompted the most basic belief for the case in question--in a manner that involves the (here assumed) obtainment of conditions (i)-(iv)--serves as the suitable basis for what flawlessly justifies that belief as well as simultaneously any and all members of the infinitely large set of ever increasingly more complex beliefs. To help show why this is so, it will be instructive to first illustrate how possessing a *positive epistemic property* clearly applies to both a most basic belief as well as even infinitely complex iterations of the relevant

sort. It seems clear, for example, that when one is a competent speaker of English, one knows (by *whatever* means) that p when and only when one knows that the sentence (in English of course--hereafter I will omit this qualification) 'p' is true. But, importantly, this would no less be the case when one replaces the first p used in the previous sentence with "the sentence 'p' is true." In other words, it also is correct that one knows that the sentence 'p' is true when and only when one knows that the sentence about a sentence "the sentence 'p' is true" is true. And so, if we grant that knowledge is closed under known entailment,³⁴ this allows for a kind of *unproblematic* complexity iteration to occur, which results in being committed to the claim that, for example, one knows that p when and only when one knows that:

The following italicized sentence that is about the bolded sentence that is itself about the underlined sentence that is itself about the double underlined sentence '*the sentence "the sentence 'the sentence "p" is true' is true*' is true.

The property of *being known* (by *whatever* means), in other words, transfers quite non-problematically across even infinitely complex iterations of the relevant kind. Put in a simpler way, it seems clear that we know that p when and only when we know that it is true that it is true that it is true... (*ad infinitum*) that p.

This shows at the very least that we should not, in general, recoil *merely* from the fact that complexity iterations like this may seem forced upon us by our professed position. And since it seems clear that being known does indeed transfer quite non-problematically to even the most infinitely complex iterations of the above kind, it should seem at least less strange at this point to claim that being known *by means of one's sense experiences* (as a result of the fact that conditions (i)-(iv) are known to obtain) does so as well for the iterations made evident by (1) and (2). Indeed, the main difference between

iterations of the sort dealt with in the previous paragraph and those described by (1) and (2) is that rather than saying we are certain, for example, simply that it is true that it is true that is true...that p; we are saying that we are certain, in effect, that p has its truth conditions *met at the moment in one way rather than another*, and that the claim that p is true has its truth conditions met at the moment *as a result of the particular way that p has its truth conditions met*, as does the claim that it is true that it is true that p, and so on ad infinitum. In our original case, for example, the truth conditions for the claim that I'm having a red sensation are *asserted outright* to be met because I am staring at a red patch of fabric while not, say, a ripe tomato. But, of course, what grounds our confidence in the *particular* way that the claim in question's truth conditions are met is our sense experiences at that time. How do our sense experiences ground things in this way? They do so (given that we are assuming that we know (i)-(iv) to hold) by allowing us to *refer* to the precise way in which the truth conditions for the claim in question are met, in a manner that is *sensitive* to the obtainment/lack of obtainment of those truth conditions. When the above conditions are met, my sense experiences allow me to refer to the truth conditions for the claim that I'm having a red sensation *only* (issues pertaining to vagueness notwithstanding) during the times when I'm actually having that kind of sensation. And, importantly, it is this *ability* to sensitively refer to the particular manner in which the truth conditions for p are met which also adequately grounds a claim like (1) as well. I would not be willing to assert (1) outright, after all, unless I was having the sense experiences which allowed me to sensitively refer to the truth conditions for my belief that I'm having a red sensation. Assertion (1) is just as much an occasion sentence as is the claim that I'm having a red sensation. I would assent to (1) *when and only when* I

would assent to the claim that I'm having a red sensation. The epistemic flawlessness of my belief that I'm having a red sensation does, for the pessimistic fallibilist, hinge on my belief that (1) being epistemically flawless. But for the just mentioned reasons I think that it is clear enough that if (granting that conditions (i)-(iv) are also known to hold) I wish to discursively set out *what my conclusive evidence is* for this last more complex belief, a completely satisfactory answer seems to be "The sense-experiences I'm now having." What makes this answer, in this situation, suitable for answering the relevant sort of skeptic *here* is clearly i) the way things are at that time; ii) the fact that it was the way things are at that time that are known to be what prompted the initial belief; and more particularly iii) that I am able to *sensitively refer*, when I actually answer the challenge, to the way things are. In such cases I can come to be absolutely certain that a particular property of a mentioned sentence (its truth conditions being met in one way rather than another) holds *as a result of*, in part, what I know my sense experiences allow me to sensitively refer to at that time.

The same holds for the even more complex assertion: *The truth conditions for "the truth conditions for 'I'm having a red sensation' are met at the moment" are met at the moment.* By being implicitly willing to assert (1), I contextually imply that I am willing to assert this last italicized claim. But we should note once again that the italicized claim is easily seen as logically equivalent to both the belief that I'm having a red sensation as well as (1). As long as knowledge is closed under known entailment, this means that I would be flawlessly justified in maintaining it when and only when I was flawlessly justified in maintaining (1) as well as a belief that I'm having a red sensation. For this reason the just-discussed properties that were shared between (1) and the belief

that I'm having a red sensation regarding what grounds each claim also holds for the relationship between this last italicized belief and, for example, (1). Given the transparency of the logical equivalence of the relevant complexity iterations it should be clear how this generalizes to even the most infinitely complex assertions of the relevant sort.

Of course, I don't expect the position that, for example, (1) and the claim that I am having a red sensation are logically equivalent to *surprise* the reader. Nevertheless, the fact that defending the pessimistic fallibilist's view on this matter depends on noting truisms like it only serves to reinforce the plausibility of the view. The interesting aspect of the account, I take it, lies in how we can use such truisms to consequently easily explain why higher-order justification of the relevant sort--even for the most infinitely complex iterations--does in fact obtain. The reason why we can easily say that we can meet the higher-order requirement of being certain, for example that it is true at a certain time that is true at a certain time that p, is precisely because it is true as certain time that it is true at a certain time that is true at a certain time that p is quite *obviously* logically equivalent to p when p is an occasion sentence. It is, in the end, precisely because it rests, in part, on such "uninteresting" truisms that the pessimistic fallibilist's view (interestingly enough!) avoids the analogous problems that the strong access internalist may face.

Let me now deal with an important objection. It might be thought that I've ignored the role that knowledge of a sentence's *meaning* plays in having it that the truth conditions for the more complex beliefs are understood as guaranteed to be met. In other words, it is obviously not *just* the fact that it was prompted by the more or less ephemeral, contingent way the world is at a certain time which makes it so that my belief

that the truth-conditions for a certain belief have been met at a certain time is itself true. What also is involved in the truth conditions of such beliefs is the meaning of the relevant sentence. If the meaning of the sentence “I am having a red sensation” were importantly different, then under the relevant circumstances that sentence would *not* be true. The meaning of the phrase therefore contributes to the truth conditions of statements that mention the phrase in question, and claim that its truth conditions are met at a certain time. In contrast, it is merely the way that more or less ephemeral, non-linguistic matters stand at a certain time which makes it so that the truth conditions for the belief that I am having a red sensation are met. Thus, what verifies these two claims outright would seem to have to be different, contrary to what I have suggested above. In order to verify that the truth conditions for the English sentence “I’m having a red sensation” are met, I obviously would have to know *both* about the way the world is at that time as well as what the phrase “I’m having a red sensation” means. It is thus at least somewhat difficult to see how the mere known fact that I can sensitively refer to the relevant nonlinguistic truth conditions could conclusively underwrite both beliefs at a certain time, given this difference in content. Obviously, the same problem would pertain to all of the more complex iterations as well.

But the claim was never that what one's sensory experiences allowed one to sensitively refer to at the time in question was the *truth conditions* for all of the beliefs, no matter how complex. They do so, given the relevant assumptions, for the most basic belief; however, as the above makes clear, that is it. The claim that was made, instead, was only that one's known ability to sensitively refer to the relevant nonlinguistic truth conditions serves *to adequately ground the flawless justification* of all of those beliefs--

ultimately, as a result of the easily grasped logical equivalence of the relevant beliefs-- and that for this reason the complexity iterations present no problem. That one must know how to speak the language in which one makes one's assertions is patently obvious. Part of knowing how to speak one's language, in turn, involves some grasp (made evident in one's linguistic and non-linguistic behavior) of the truth conditions of simple enough assertions. But this does *not* mean that being implicitly willing to assert outright that the truth conditions for a simpler assertion are met at a certain time requires examining at that same time whether one does in fact have a grasp of what the truth conditions for the statement in question in fact are. One's implicit willingness to make the higher-order assertions, rather, merely serves as what follows from the momentum of the original, simplest assertion. Recall that (via condition (i)) it is being merely *assumed* in this case that no skeptic exists who is concerned at all with quibbling over our grasp of our mother tongue. Moreover, in reflexive instances of skeptical questioning, quibbling with oneself on this matter seems at best idle and at worst bound to needlessly confuse. The above discussion has been aimed at addressing what might look like a theoretical *problem* with pessimistic fallibilism. To that end, what I have tried to show is that when I know that I'm able to sensitively refer to the truth conditions for the most basic belief, I have thereby positioned myself so as to be absolutely certain of any of the more complex beliefs generated by the complexity regress as well.

Now, there is obviously plenty of room for various skeptics to voice doubts over whether conditions (i)-(iv) in fact obtain. Any such skeptic must be dealt with in order to establish that a certain belief is epistemically flawless *in those contexts*. Doing so will no doubt require that one go on to flesh out *what grounds one has* for claiming that one's

sense experiences allow one to sensitively refer to the truth conditions for the most basic assertion. One skeptical of another's ability to employ their sense-experiences to sensitively refer to certain kinds of truth conditions would have to find fault with any argument so offered. But my *present* goal has been achieved as long as it is clear that *when* one is certain that one's sense experiences allow one to sensitively refer to the truth conditions for the most basic belief, the basic belief as well as all of the complex ones generated by the complexity regress are thereby proven to be flawlessly justified. Debating the extent of the power that one's sense experiences have for getting one into sensitive contact with various sorts of truth conditions is a relevant, but quite separate matter. In this we see some of the flexibility involved in the pessimistic fallibilist's general epistemological outlook. I will, later on, make use of what this sort of flexibility affords.

2.3 Conclusion

This concludes my examination and development of some of the relevant aspects of pessimistic fallibilism. By framing an epistemological approach around the fact that convincing skeptical arguments can be formulated which challenge the assertions we all are willing to make about the external world, the pessimistic fallibilist stresses the importance of backing up higher-order commitments for grounding these assertions. We have found that this commits the pessimistic fallibilist to claiming that properly backed up assertions require the existence of an infinitely large hierarchy of ever-more-complicated flawlessly justified beliefs. I have then tried to show why--in contrast, perhaps, with the related regress of the strong access internalist--this presents no problem for the pessimistic fallibilist.

As the above no doubt makes clear, I believe that when looked at generically enough, the pessimistic fallibilist's approach is basically sound. Having one's views regarding certain perennial philosophical issues grow primarily out of a theoretically-neutral search for what sorts of beliefs survive skeptical attack in various contexts, is a healthy and paradigmatically philosophical intellectual exercise. Where the pessimistic fallibilist goes wrong, in my view, is in his implicitly maintained view that perceptual experience is only strong enough to put us into sensitive "contact" with the truth conditions for experiential beliefs. But here I am again getting ahead of myself. Before explaining what I find wrong with pessimistic fallibilism, I would like to round off the more expository part of the overall work by explaining what I find wrong with optimistic fallibilism of the non-discursive brand. It is towards this matter that I will now turn.

CHAPTER III NON-DISCURSIVE OPTIMISM

3.1 Introduction

Initially, I defined a non-discursive optimist as anyone who believes that our inability to refute a skeptic regarding the epistemic flawlessness of sensory-experience-based statements about the external world is compatible with our nonetheless being epistemically justified in maintaining those same sorts of claims. I now turn to how a view such as this fares as a response to the pessimistic fallibilist--here understood as the claim that fallibilism entails that believing in the (likely) truth of external world claims is presumptuous at best. I will critically examine the way that non-discursive optimism fares by, in particular, examining the views of the later Wittgenstein, contextualists, process-reliabilists, and naturalized epistemologists. The discussion will presuppose what has been discussed and developed so far, and thus certain substantive conclusions reached in that discussion are going to do a lot of the work. Let me now both briefly recapitulate and slightly develop some of the important points.

First, and perhaps most importantly, the pessimistic fallibilist makes no claim to have grasped and made use of an *intuitive* understanding of "our concept" of knowledge and/or justification. Instead, his position is designed to gain traction and relevance based upon the contingent fact that most people take themselves to be in a strong epistemic position with respect to various external world claims. The problem of philosophical skepticism arises because this is in clear conflict with following claim: *since fallibilism holds it follows that to claim that any external world belief is (likely to be) true is presumptuous at best.* In other words, according to the pessimistic fallibilist, fallibilism entails that we must consider external world claims as no more worthy of assertion than

any other claim that we consider merely presumptuous (hopes, hearsay, wild sounding rumors).

Importantly, the pessimistic fallibilist takes his position to arise fundamentally from a desire to remain neutral on otherwise controversial theoretical matters. Aside from the above results, we have seen that other interesting things fall out of this desire for neutrality. For one, the skeptic is not required to claim that one can arrive at his view by employing a form of *a priori* reasoning that results in known truths distinct in kind from contingently true claims. For all the pessimistic fallibilist knows, there is no such thing as *a priori* knowledge. Relatedly, the pessimistic fallibilist is not committed to there being any necessary truths distinct in kind from contingent truths.

But, are there other substantive meta-epistemological matters that the pessimistic fallibilist must take a stand on? Are there matters that the non-discursive optimist in particular will wish to point out are indeed controversial? As already discussed, the key issue according to the non-discursive optimist pertains to what is supposed to follow from the fact that once we grant that our sensory evidence is consistent with the truth of the skeptical scenarios, it follows that believing in the (likely) truth of anything about the external world is presumptuous at best. Given the discussion of the first chapter, that this indeed follows is, I take it, clear enough. But, of course, the *non-discursive* optimist will wish to point out that being "presumptuous at best" in the particular sense at hand is quite different from being *in fact* unjustified. It may be true that our sensory evidence is consistent with the skeptical scenarios, and consequently that in the context of attempting to *argue* for the (likely) truth of external world claims we will be forced to admit *to some challenger* that we are being presumptuous at best in maintaining our sought-after

conclusion. But, the non-discursive optimist asks, why does the reasonability of external world claims depends on being able to perform the fool's errand of refuting that challenger? This is, I think, the most natural cool-headed reaction to the pessimistic fallibilist's position as I have presented it. Most people will, I think, readily grant that we don't *know with complete certainty* that we are not dreaming, and that therefore our epistemic position is weak in that we can't refute a skeptic in an actual debate. But these same folks will also deny that we *should* have to--where the sense of "should" here pertains to what being *in fact* epistemically justified requires. They will deal with the admitted force of the arguments of the first chapter that seem to fall out of realizing the full implications of our fallible epistemic predicament, then, by simply divorcing someone's being in fact epistemically justified who is in that epistemic predicament from that same someone's ability to refute a skeptic. Rather than accepting skepticism, they might then try to come up with convincing arguments of their own for why being in fact justified has nothing to do with being in a fallible epistemic predicament and also being able to discursively spar with some wily challenger.

Can the non-discursive optimist show thereby how the pessimistic fallibilist has gone out on a limb? Is defending non-discursive optimism the way out of the difficulty? To address these questions we must get clear on what it is to be in fact justified, and, more importantly, what it means to reasonably address--as a non-discursive optimist must--the very question of what humans are in fact justified in believing.

One way the relevant matter can be broached is by drawing a distinction between being *propositionally justified* versus *doxastically justified*. Propositional justification, it is claimed, pertains to what we are *in fact justified* in believing. One is propositionally

justified in believing something when one in fact has justification for that claim. But, importantly, one can be propositionally justified in believing something, on this view, without actually explicitly believing it. One may in fact have strong evidence for the claim that one is not a French waiter, for example, without ever explicitly believing that one is not a French waiter. Doxastic justification, in contrast, pertains to beliefs one does explicitly hold. One is doxastically justified in believing that *p*, roughly put, when one is propositionally justified in believing that *p*, one actually believes that *p*, *and* where the fact that one is propositionally justified (or the facts that entail that one is so justified) played a central enough role in the actual formation of the belief that *p*.³⁵

Questions about whether or not someone is propositionally justified pertain to what is also sometimes called *the structure of justification*. Those who hold a view about the structure of justification wish to explain what we are in fact justified in believing by attacking the matter at an extremely general level. Two different views about the structure of justification, for example, are that of the foundationalist versus coherentist. The foundationalist believes that generally speaking there are two kinds of justified belief: *inferentially justified beliefs* and *non-inferentially justified beliefs*. Inferentially justified beliefs are those that are justified *solely* in virtue of the fact that they are able to be supported by an inference from a different justified belief, where the latter belief is understood to be *epistemologically prior* to the inferentially justified belief. A belief that *p* is said to be epistemologically prior to another belief that *q* when and only when the justification of *q* depends on the justification of *p* (in virtue of their inferential relationship), but not vice versa. So, inferentially justified beliefs are *propositionally* justified when and only when a subject holds epistemologically prior justified beliefs.

The same inferentially justified beliefs are doxastically justified when and only when a subject forms them in a way that makes the epistemologically prior belief play a central enough role. Non-inferentially justified beliefs, in contrast, are beliefs the propositional justification of which does not depend on any epistemologically prior beliefs. In the case of beliefs about contingent reality, the foundationalist will claim that non-inferentially justified beliefs are propositionally justified *solely* by the fact that one has certain sorts of sense experiences that have the right kind of relationship to the kind of belief in question (where what it means to "have the right kind of relationship" here can mean quite different things depending upon the brand of foundationalism espoused). The same beliefs are doxastically justified when and only when one forms them in a way that makes one's sense experiences play a central enough role.³⁶

There are well-known theoretical problems with foundationalism.³⁷ The most interesting criticisms come from the coherentist. Coherentists deny, in effect, that there are non-inferentially justified beliefs. It is not necessary to here address the various arguments that coherentists have offered (see the works cited n. 37). Instead, here we need only get a rough feel for the way in which their view differs from the foundationalist regarding the structure of justification. Unlike foundationalists, coherentists typically claim that particular beliefs are justified *holistically*.³⁸ This means that a belief is propositionally justified in virtue of its coherent relationship to either all of the other beliefs that a subject holds, or else a good (relevant) portion of that entire belief set. The relationship of coherence can be spelled out in various ways. To give just one example, a belief might be claimed to cohere well with one's relevant set of beliefs when the truth of that set of beliefs *makes objectively probable* the belief in question. Note how this differs

starkly, of course, from the foundationalist's claim that a belief can be propositionally justified in virtue of its relationship to one's sense experiences alone.

The above discussion of the difference between foundationalism and coherentism will not be extremely illuminating to one who is not already familiar with the general character of this debate. But I mention them at this point only to give the reader a rough feel for what it can mean to have a certain view about the structure of justification. The important issue to now address is whether arguments for a certain view about the structure of justification (and hence about what it means to be propositionally justified) can be deployed to ameliorate the unpleasantness of (if not effectively counter) what follows once we grant fallibilism.

Well, clearly enough there would be nothing problematic about what follows from fallibilism, *if* we knew independently that the structure of justification was, in fact, arranged so that, say, *our own* perceptually-based beliefs about the external world were fallibly justified anyway. So, arguing in the general way that the non-discursive optimist does, if nothing else, introduces an interesting and relevant complication into the fold. But what does it *mean* to figure out what the structure of justification is in our own case? How is it that we could go about addressing that very question, specifically regarding the justificatory status of external world beliefs? It seems that we have, generally speaking, two viable options. We can go about things by granting fallibilism and then asking which of our external world beliefs that we already possess are nonetheless made probable enough by our evidence for them, and which are not--where we also leave the relevant matter open at the outset of our investigation, for *all* of our external world beliefs. Contrarily, we can ask what the structure of justification is for human beings in general

granted that human beings possess the fallible but still reliable faculties that we already take them to possess (or else granted that the methods of investigation we use to discover the way humans are constructed are able to deliver us with knowledge of that matter) and then apply what we discover about what it means for such creatures to be justified to ourselves.

Certain non-discursive optimists—among them process-reliabilists and naturalized epistemologists--pursue the second option. In order to divorce their own professed fallible predicament from being "really justified" *themselves*, then, it seems clear that a non-discursive optimist of this sort must *first* understand themselves to be a being with a certain set of faculties--perhaps, for example, faculties that allow for sensitive and reliable perceptual contact with external reality--and *then* point out why the possession of *such faculties* is both what constitutes being justified as well as something that is compatible with the failure to be able to perform the relevant argumentative maneuvers. A popular non-discursive optimist maneuver is to try to establish the point about what it means to be justified by constructing thought experiments in order to mine for intuitions about whether or not subjects are justified in various cooked up scenarios. Non-discursive optimism is thus, in this case, wedded to the idea that we can discover what it means to be justified by figuring out what we are disposed to say about various subjects in various situations involving the formation of external world beliefs. By examining cases, on this view, we can get the account of justification *right* and thereby see that being justified has got nothing really to do with being able to refute a skeptic. From there it only takes the raising of considerations for why one is just like the subject involved in the cases where the belief is justified to explain why one's own ability to

refute a skeptic has got nothing to do with one's own beliefs being justified. (And, of course, the same philosopher would not be bothered to learn that they are unable to refute someone who denies that they are like the subjects in the favorable cases.)

Now, obviously the pessimistic fallibilist will have none of this. What does it mean, for example, to employ thought experiments to get at what it means to be "really" justified? What does it mean to get an account of a concept right? These are highly theoretical meta-philosophical questions. It is not in any way clear what answer they have--nor, more importantly, is it immediately clear what it would even be to go about answering them in an illuminating rather than merely *self-reassuring* way--and thus at best it seems obvious that the non-discursive optimist has a lot more arguing to do at the relevant meta-level before they can put stock in their general strategy. And yet, here I would like to urge the following point: for basic *epistemological* reasons it is doubtful that any argument they do offer on this matter will be helpful. It should be quite clear that the methodology of uncovering the structure of justification by the examination of thought experiments *might*, for example, confirm nothing whatsoever about what it means to be "really" justified. What if, after all, one's intuitions (or the intuitions of one's linguistic community) fail to track the relevant truth here? Again, much more importantly, what would it even *mean* to track the relevant truth here? *Basic epistemology* should force the non-discursive optimist to see that to claim that their results are "intuitive" will do *nothing* to support the claim that they track the "conceptual truth" at hand, unless, of course, they have an understanding of why knowledge of a claim's intuitiveness (for them or their linguistic community) makes the relevant conceptual truth likely (or guaranteed). This skepticism over the results of conceptual analysis, of course,

falls directly out of the pessimistic fallibilist's general epistemological outlook. When we mine a thought experiment for own intuitions, all we know with certainty is that we find, for example, the claim "S is justified (or not)" *intuitive*. What we lack is knowledge of what being found intuitive has got to do with confirming what it is to be "really" justified. Armchair conceptual analysis may, for all any of us know, be without any objective import *whatsoever*. And, ironically enough, armchair reflection reveals this.

Nevertheless, it is equally true that in a qualified way the pessimistic fallibilist does wish, like the non-discursive optimist, to inquire about what can accurately enough be called the structure of justification in their own case regarding external world beliefs. But they pursue the first option mentioned above, while--here is the important qualification--making no claims to have thereby uncovered what it is to be "really" justified *generally speaking*. They pursue a methodology that asks, among the beliefs about the external world that they already possess, what their status is vis-à-vis the evidence they possess. They inquire after what might as well be called a kind of *fact*, then, regarding their own epistemological situation, while leaving open the question of whether or not the facts thereby uncovered correspond to being "really" justified. In this last sense, of course, they differ from the non-discursive optimist. They also differ, more importantly, in that they do not assume at the outset that any of their beliefs in the relevant domain have a particular epistemological status. As such, they will *not* take it as granted that human beings are built a certain way (or that they possess a reliable means for discovering how human beings are built).

Here we arrive at the important point: when one proceeds in a manner that grants fallibilism and then questions the epistemological facts about one's own external world

beliefs in a manner that starts from a neutral, first person standpoint regarding that matter, it seems clear that the question of what such facts are becomes identifiable with the question of which of one's beliefs can survive (various kinds of) challenges. To state the point differently: when I grant fallibilism as well as *only* that I have certain external world beliefs, and attempt to go on to discover the strength of my own epistemological situation regarding such beliefs, the *only* fruitful methodology will be one that involves determining which of my beliefs survive (various kinds of) skeptical attack. The pessimistic fallibilist's reason for letting all matters of "epistemological fact" rest here on the question of whether or not any of their external world beliefs survive skeptical attack, then, is nothing other than that this is the only way to uncover substantive results when one also proceeds in a manner that fails to assume at the outset that any of one's external world beliefs are or are not already made objectively probable by one's evidence.

Most importantly of all, when proceeding as such, it does look to be impossible to divorce being shown to be presumptuous at best from being in an *extremely weak* epistemological position. Again, however, it needs to be emphasized that the pessimistic fallibilist makes no claim regarding whether or not this position does or does not correspond to being "really" justified. If pressed on the matter, he will back off from calling his results "epistemological facts," and merely say that they are conclusions he reaches *via a neutral, first-person methodology*.

This, of course, invites the question: *why* proceed with that sort of methodology? Why proceed in a manner that fails to assume at the outset any epistemological facts about the beliefs under examination? Here the ready answer seems to be that doing so is supposed to satisfy a peculiarly philosophical kind of curiosity.³⁹ But the pessimistic

fallibilist has a stronger answer. As long as one is someone who is willing to assert things about the external world, as we have seen, this implies that one (at least implicitly) takes oneself to have *some positive evidence* for such claims in the sense that one takes oneself to be positioned to show someone who doubts the (likely) truth of an external world belief why they are wrong. In this way, one takes one's epistemic position regarding such claims to be strong; or at least stronger than claims like that the population of New York at the moment is even, or that $2+2 = 5$. By being willing to assert the relevant claims, one takes the "epistemological facts" to be favorable in this way. This means that one is, as a matter of contingent fact, readily prepared to dig oneself into the relevant epistemological hole. Now, importantly, *I do not see why this willingness to assert things about the external world would go away throughout the duration of a neutral, first-person examination of the epistemological status of one's beliefs* (though I will hold off adequately explaining why this is so until the next chapter). But, also true is that when we do so (we are still assuming) fallibilism looks plausible, and can be seen to entail that those same claims are presumptuous at best. As a result, all *fallibilists* will have to, it seems, admit that they will be unable to back up the assertions that they also admit they are continuously willing to make. *We should* seek to discover if we can back up such assertions starting from neutral ground, then--in other words, we should discover if we can satisfy the kind of philosophical curiosity alluded to above--only because our inability to do so in a way that *proves* that we have some evidence for our external world claims will result in *a paradox*. The paradox arises because we both believe we have some basis for making external world claims (given that we all are willing to make the relevant assertions) even though we accept the plausibility of fallibilism and hence must

accept the position that those same claims are presumptuous at best. This is the paradox. It completely describes the *problem of philosophical skepticism* according to the pessimistic fallibilist.

Of course, there may as a contingent fact be people who are *not* willing to make the relevant assertions about the external world. But for these folks the pessimistic fallibilist's negative arguments will not be of any interest. His view would be akin to skepticism over whether the population of New York City at the moment is even, or over whether $2 + 2 = 5$. But to say that skepticism over external world claims has, unlike these last claims, garnered interest is, of course, a severe understatement. Thus, I take it as obvious that we all take our epistemic position to be such that we have at least *some* positive basis for external world claims. We *must* pursue the question of whether we can prove via sound, non-question begging argument that some of our beliefs about external reality are in fact properly evidentially backed, then, for no other reason than that our failure to be able to do so results in paradox. Here I would like to stress how the fact that this kind of paradox exists is quite separable from any claim that we are "really" unjustified in our external world claims. It is true that what we are after is, in a way, perhaps accurately enough describable as an answer to the question of what the structure of justification is for any of our external world beliefs; i.e., what the "epistemological facts" are for such beliefs. But the crucial qualification is that we understand this investigation to be an uncovering of facts about the amount of evidence we possess for our external world claims *according to one way of going about answering that question--* a way which is quite different from how the non-discursive optimist goes about matters. Via *this* kind of examination a paradox is born. Nor, as we have seen, can we (without

additional argument) plausibly divorce being "really" justified from being able to satisfy the relevant kind of philosophical curiosity. Any attempt to separate matters in this way will, as I have pointed out above, run into all sorts of other forceful, legitimate skeptical worries regarding one's ability to independently discover what it is to be "really" justified. So, the upshot is that as long as we are willing to assert things about the external world and we are able to see the futility in discovering what it is to be "really" justified in the just-mentioned sense, we should be concerned with satisfying the relevant kind of philosophical curiosity. Otherwise, we would have to live with the paradox, and this is simply not an option for a rational, reflective thinker.

This presents the proper backdrop for the way in which I wish to respond to the non-discursive optimist's general strategy for dealing with skepticism. Let me now deal specifically with certain proposals.

3.2 Wittgenstein

In his book *On Certainty* Wittgenstein expresses his (evidently) considered views regarding the skeptical problem.⁴⁰ Wittgenstein believes that the claims that are often subject to skeptical attack are, in most ordinary contexts, "propositions [that] are exempt from doubt, [that] are as it were like hinges on which...[doubts] turn" (par. 341). Such "hinge" propositions are statements that we would indeed assent to if, for example, we were questioned on the matter. But most people would, for any such proposition, also find the question regarding what he/she thought on the matter *extremely odd*. This is because acceptance of hinge propositions results, according to Wittgenstein, *not* from the fact that we were ever taught them (as children, say), but instead from the fact that they help to describe the fundamental aspects of a typical socially indoctrinated "frame of

reference" (par. 83), "picture of the world" (par. 94); or, perhaps most accurately, *way of life* (pars. 344, 358). These are the sort of propositions that, usually, are not explicitly considered, but nevertheless are such that our (often) implicit acceptance of them is what helps to cause our *practices* of confirmation/disconfirmation--i.e., our ordinary methods for deciding upon whether we know or do not know that some mundane matter holds--to proceed as they in fact do (pars. 95-98, 105, 151). The following passages help to develop the pertinent idea:

If I say "*we assume* that the earth has existed for many years past" (or something similar), then of course it sounds strange that we should *assume* such a thing. But in the entire system of our language-games it belongs to the foundations. The assumption, one might say, forms the basis of action, and therefore, naturally, of thought (par. 411).

Giving grounds, however, justifying the evidence, comes to an end;--but the end is not certain propositions' striking us immediately as true, i.e., it is not a kind of *seeing* on our part; it is our *acting*, which lies at the bottom of the language-game (par. 204).

The existence of the earth is rather part of the whole *picture* which forms the starting-point of belief for me (par. 209).

Now it gives our way of looking at things, and our researches, their form. Perhaps it was once disputed. But perhaps, for unthinkable ages, it has belonged to the *scaffolding* of our thoughts...(par. 211).

...I did not get my picture of the world by satisfying myself of its correctness; nor do I have it because I am satisfied of its correctness. No: it is the inherited background against which I distinguish between true and false.

The propositions describing this world-picture might be part of a kind of mythology. And their role is like that of rules of a game; and the game can be learned purely practically, without learning any explicit rules (pars. 94-95).

If someone doubted whether the earth had existed 100 years ago, I should not understand, for *this* reason: I would not know what such a person would still allow to be counted as evidence and what not (par. 231).

That I have hands, for example, is a proposition whose implicit acceptance on my part is *exhibited behaviorally* by, e.g., how I would use my hands to open the fridge door to verify/disconfirm whether or not there is a soda in the fridge, and by how I might (more or less explicitly) verbally convey to someone that *that's what I did* if they asked how I knew that there was soda in the fridge. In like manner, that the earth is more than a few hundred years old is something whose acceptance on the part of the historian is exhibited by their inquiries into what happened to so and so in *antiquity*. In this way, hinge propositions are best seen as *non-epistemically* related to the subject who employs them. That this is what Wittgenstein maintained is confirmed by the following passage: "At the foundation of well-founded belief lies belief that is not founded" (par. 253). In other words, normally we would not assent to the statement "I have 5 fingers" because we think we have *evidence* for it. The statements "I have hands", "I have 5 fingers", etc. are usually treated by us as just as certain as anything we could offer in support of their truth (par. 250). To repeat their relevant feature, without statements like them in place, we would not undergo the specific, socially imbued *practices we in fact implement* in coming to doubt/believe anything whatsoever.

We check the story of Napoleon, but not whether all the reports about him are based on sense-deception, forgery and the like. For whenever we test anything, we are already presupposing something that is not tested...(par. 163)

Think of chemical investigations. Lavoisier makes experiments with substances in his laboratory and now he concludes that this and that takes place when there is burning. He does not say that it might happen otherwise another time. He has got hold of a definite world-picture--not of course one that he invented: he learned it as a child. I say world-picture and not hypothesis, because it is the matter-of-course foundation for his research and as such also goes unmentioned (par 167).

Obviously, then, on this view I do *not* assent to such a statement like "I have hands" because I take it to be *a proposition that I have evidence for*. And thus, contra to what I

have suggested again and again, according to Wittgenstein I do not assent to such propositions because I'm (even implicitly) willing to *assert* hinge propositions. Placed as it is within my epistemological practices, a statement like "I have hands" is not, *as such*, what the further practices of confirmation/disconfirmation I usually employ could, e.g., show to be false. Rather, my acceptance of it and statements like it determines how those very practices occur. And *of course* it would be a mistake of logic to maintain that hinge propositions are susceptible to falsification/verification *via* the very methods (implicitly) accepting such propositions helps to determine.

The most important aspect of this view, I think, is the claim that the existence of *actually* falsified/verified claims depends on the existence of accepted hinge propositions; which are a type of non-falsified/non-verified claim. The above discussion should give the reader at least a workable feel for the proper way to construe this dependence. Let us paraphrase the point as follows: *epistemological beliefs require non-epistemological beliefs*.

Clearly, then, Wittgenstein espouses a kind of non-discursive optimism. We are justified in believing what we do, according to Wittgenstein, because to be justified involves socially inheriting a certain way of acting. To be justified in believing something, then, involves nothing more than behaving (both verbally and non-verbally) in accordance with one's inherited world-picture. Whether or not one is "really" justified will depend on whether one acts, or fails to act, in accordance with the associated established practices. Importantly, according to Wittgenstein, for the already discussed reasons refuting someone who challenges any *hinge proposition* cannot, given the logic of inquiry so employed, be part of that practice.

One might simply say "O, rubbish!" to someone who wanted to make objections to the propositions that are beyond doubt. That is, not reply to him but admonish him (par. 495).

If someone wanted to arouse doubts in me and spoke like this: here your memory is deceiving you, there you've been taken in, there again you have not been thorough enough in satisfying yourself, etc., and if I did not allow myself to be shaken but kept to my certainty--then my doing so cannot be wrong, even if only because this is just what defines a game (par. 497).

The queer thing is that even though I find it quite correct for someone to say "Rubbish!" and so brush aside the attempt to confuse them with doubts at the bedrock,--nevertheless, I hold it to be incorrect if he seeks to defend himself (using, e.g., the words "I know") (par. 498).

A natural reaction here, however, runs as follows: it is true enough that statements like "I have hands" are not *normally* the sort of statement we might seek to verify/falsify. But this alone does not show that the activity of doing so--that is, seeking to verify/falsify what are normally hinge propositions--is somehow misguided (Williams 1996, 27). Asking whether we know *anything at all*, for that matter, seems to be nothing more than a kind of highly generalized "anthropology" (Stroud 2000, 122-123). In this sense, the kind of questions asked by philosophers, although indeed unusual, may also be no more or less confused than any unusual question that a cultural anthropologist, for example, might ask.

As I see it, however, the Wittgensteinian response to this reaction is forceful and simple. It also describes what I think Wittgenstein gets completely right. Namely, to say that epistemological beliefs require non-epistemological beliefs is to make an *accurate* claim about the logic of inquiry. As a result, although we can indeed make what were formerly non-epistemological beliefs subsequently epistemological beliefs; the philosopher goes wrong in thinking that the discoveries they make regarding the status of what are, in normal contexts, hinge propositions are in any way relevant for determining

the status of epistemological beliefs in those same normal contexts. Of course it may be that we find that a hinge proposition like "I have hands" is poorly verified *when we make the relevant matter a subject for inquiry* rather than allow the claim in question to remain a hinge proposition. But the whole force of the Wittgensteinian position can dawn on one once it is realized that issues between two systems of inquiry that differ in terms of their accepted hinge propositions--what we might call issues between two different "world pictures"--are not (possibly?) resolvable by epistemological means; *and* that this fact poses no epistemological problem for either world picture given how the logic of inquiry really works. Thus, all we might be able to do to get a philosopher with a completely different, and maturely thought out world-picture to come around to viewing things in a more "common sense" manner is try our best to *persuade* him/her to simply drop the odd sounding world picture that makes a proposition like "I have hands" something other than a hinge proposition. The *last* thing we should do is try to argue with the philosopher on this matter. I think that anyone who has paid even a scant amount of attention to how people actually argue should be able to notice how Wittgenstein seems quite correct on this point. An accurate epistemology must accommodate it.

If someone said to me that he doubted whether he had a body I should take him to be a half-wit. But I shouldn't know what it would mean to try to convince them that he had one. And if I had said something, and that had removed his doubt, I should not know how or why (par. 257).

I would like to reserve the expression "I know" for the cases in which it is used in normal linguistic exchange (par. 260).

I can imagine a man who had grown up in quite special circumstances and been taught that the earth came into being 50 years ago, and therefore believed this. We might instruct him: the earth has long... etc. --We should be trying to give him our picture of the world.

This would happen through a kind of *persuasion* (par. 262).

But why *am* I so certain that this is my hand? Doesn't the whole language-game rest on this kind of certainty?

Or: isn't this 'certainty' already presupposed in the language-game? Namely by virtue of the fact that one is not playing the game, or is playing it wrong, if one does not recognize objects with certainty (par. 446).

Supposing we met people who did not regard [relying on currently accepted physics] as a telling reason [for believing something]. Now, how do we imagine this? Instead of the physicist, they consult an oracle. (And for that we consider them primitive.) Is it wrong for them to consult an oracle and be guided by it?--If we call this "wrong" aren't we using our language-game as a base from which to *combat* theirs?...

I said I would 'combat' [those who disagree with me],--but wouldn't I give [them] *reasons*? Certainly; but how far do they go? At the end of reasons comes *persuasion*...(pars. 609, 612).

So Wittgenstein's view is that the kind of "anthropology" envisaged by Stroud (ibid.) and others cannot produce a substantive epistemological result for those who refuse (in a way, properly, as Wittgenstein sees matters) to subject what are hinge propositions for them to epistemological scrutiny. So, it is not so much that the activity of subjecting hinge propositions to epistemological scrutiny is misguided in Wittgenstein's view, but rather that it is misguided to think that the results of this activity have any bearing on the epistemological status of the beliefs of normal folks. The beliefs of normal folks remain justified, even though philosophers reach a different verdict when they perform their peculiar sort of examination.

But I think Wittgenstein would have to admit that there *is* a way for even normal folks to recognize why subjecting some of their hinge propositions to epistemological scrutiny wouldn't be a fool's errand. The reasons we might seek to do so are compatible with thinking it completely misguided to think that the errand in question would affect the epistemological status of normal beliefs in normal contexts. Instead, we can see the activity of the philosopher to have a different, related sort of relevance. What I have in

mind is the following simple enough realization. Suppose there are two subjects X and Y that have an otherwise (more or less) identical world picture, except for the following difference. Some of the propositions that are hinge propositions for X are rather *verified* epistemological propositions for Y. Is there a way of comparing the epistemological situation of X and Y? I would think that there is. Namely, we must say that Y is in a *stronger epistemic position* than X. This is not an odd thing to claim. The two subjects share a world (though not a world-picture), after all, and Y simply knows more about that world than X. What X takes implicitly for granted as part of his/her world picture, Y has actually gone out and verified. In like manner, we all consider the expert on a certain subject matter to be in a stronger epistemic position than a novice; and not *just* because the expert knows more facts than the novice. No, we recognize an expert to be in a stronger epistemic position because of how a novice may, for example, learn from a textbook that contains claims that a certain expert has *actually researched*. Surely an expert has a different world picture--in the relevant sense--than a novice, and yet we can *still* compare their two world pictures in the above epistemologically relevant way.

Because of this, each of us can now ask, quite legitimately: how strong can my own epistemological situation become? How many of my currently accepted hinge propositions can be verified? How many will be unable to be verified? How many will be falsified? The pessimistic fallibilist has an answer to the second question: none. For this reason, the pessimistic fallibilist can be understood, accurately enough, as claiming that a common sense epistemological position is the *strongest one possible*. Wittgenstein has said nothing inconsistent with this view; and he would *have* to agree with this position, of course, if he also agreed that we can compare world-pictures in the above-described

sense. Therefore, what he says is completely compatible with pessimistic fallibilism; we just have to phrase the central claim of the latter position differently in order to state it in terms that a Wittgensteinian will accept. But, as we have just seen, this is not very difficult.

This having been said, let us now return to the point that I think is absolutely essential for understanding the problem of philosophical skepticism: it seems that even common sense folks would find the claim that, say, an "everyday scientific" epistemological situation is the strongest one that they can achieve to be *paradoxical view*. If this is right than even normal folks can feel what, in my view, describes the full force of the problem of skepticism; albeit here put in Wittgensteinian-friendly terms. So, consider again what one takes to be the epistemological status of the following three beliefs:

(1) I have hands

(2) The number of people in New York City right now is even

(3) $2 + 2 = 5$.

In my own case, I would treat the claim that my epistemic position is, in a specific sense, weak with respect to (1) - (3) quite differently, depending upon which claim was at issue. I would take it as *obvious* that my epistemic position is not strong with respect to both (2) and (3). But, this is only because I take my epistemological position to be such that I have strong evidence for the negation of (3) and no evidence one way or the other for (2). Now, to say that my attitude with respect to (1) is similar is wildly incorrect. It just seems rather odd to claim that I am not, at the moment, in a strong epistemic position with respect to (1). Given how I *contingently* take my epistemic position to be, nothing, I

would think, *should be easier* than *showing* someone who doubted its truth that (1) holds. Therefore, if it does turn out that (1) is presumptuous at best, then this is indeed *extremely* paradoxical, which is a problem. I take it that even non-philosophical folks will agree, though perhaps Wittgenstein is right and I am completely wrong here, or else missing something subtle. When it comes to what is supposed to fall out of a scrutiny of ordinary cases, it seems one must always defer to the brilliance of Wittgenstein. But even in this case it would just be that the skeptical problem is a very *personal* one for me (which, given the relevant methodology, is a result that is not surprising at all, of course). The reader must decide for themselves whether the relevant paradox arises in their own case.

But even if it is granted in my favor that the paradox does arise for a significant number of folks (even if those folks are just philosophers), making the problem of skepticism quite relevant; there is a deeper issue at stake here that needs to now be addressed. It has to do with whether a critical examination of the semantics of knowledge claims can reveal that the *status of knowledge claims themselves substantively changes* when we oscillate from ordinary to philosophical contexts of assertion. To say that it does is, of course, quite within the general spirit of Wittgenstein's response to skepticism. If a substantive change of a certain sort can be pinpointed, then perhaps the conclusions about our knowledge that are reached in philosophical contexts can be demonstrated to be *only conditionally problematic*. Perhaps, by bringing the character of the relevant change into proper relief, we can determine that the paradox arises *only* when we do in fact subject hinge propositions to epistemological scrutiny. If so, then the way for one to avoid the paradoxical result is to simply (and quite legitimately) refuse to subject such claims to

that kind of scrutiny. Addressing this issue requires a critical examination of the contextualist's response to skepticism.

3.3 Contextualism⁴¹

Following the terminology employed by Peter Unger (1984) and others, one can distinguish two opposing views surrounding the semantics of *S knows that p*: *invariantism* versus *contextualism*. Contextualists hold that when a speaker makes a claim to everyday sorts of knowledge such as:

- (1) I know that Barack Obama is currently President of the United States

what he *means* is paraphrasable as something like:

- (2) My epistemic position with respect to the claim that Barack Obama is currently President of the United States is strong enough to meet the current conversationally-determined standards for knowledge.

So, it is the contextualist's claim that one can know things like the claim found in (1), as long as (and here is the crucial point) the standards for certainty explicitly or implicitly in place for the utterance are lax enough. It is the lax enough standards which will allow a particular knowledge claim like (1) to be considered true when the claim found within it in fact meets such standards.

The view springs from a more general semantical view that is Wittgensteinian in spirit. It is that assertion meaning is inextricably dependent, at least in part, on the larger context surrounding the particular utterance made. According to the contextualist, what is a true assertion in one context need not be true in another context, and this is so even when the same putative referent of, for example, any singular term that might be found within the assertion is involved in both assertions, and where the identical expression is, for example, audibly uttered. So, on this view the assertion

(3) That field is flat

uttered under the appropriate circumstances, might be true in a context where a group of people are looking for a nice place to picnic, but false in a different context where a group of people, staring at the very same open expanse (the same referent of "that field") are looking to play a game of croquet. In particular, the contextualist will claim that what changes the truth value here surrounds the semantics of "is flat" when uttered in (3).

The term "flat" is a so-called absolute term. A thing is flat when and only when it is absolutely or really flat, which is to say, without any bumps or perturbations whatsoever. So, if a contextualist wishes to say that an utterance involving an absolute term nonetheless expresses something true in one context, but false in another, they are required to make sense of an assertion like (3) in a way that, obviously, does not make its semantics identical in each case. Even though in each case it is (3) which is audibly uttered, and it is the same field in question that is the referent of the singular term involved in the utterance, the contextualist will claim that what is meant by the words themselves in, say, the picnicking case is something like

(4) That field is flat enough to make for a nice picnic.

In other words, to make the Wittgensteinian parallel clear, the contextualist will claim that an utterance like (3) is *used* in the relevant context imagined as a means for getting the party to notice the area in question as one which will suit their purposes; and this particular sort of *use* fixes its meaning. It is *not*, on this view, meant as something like:

(5) That field has no bumps or perturbations whatsoever.

According to a contextualist, one is wrong to take the utterance of (3) in the context in question as, even strictly speaking, meaning something like (5).

It is by making much of the distinction between what the utterance *qua utterance* means versus how a speaker *uses* an utterance, however, that a so-called *invariantist* will find room to protest. The invariantist will claim that utterances mean what they do regardless of what the context happens to be. So, on this view an utterance like (3) always *means* something like (5), no matter the context. Of course, the invariantist will admit that in the above-described context the speaker is *using* (3) for the purposes which are indeed captured well enough by (4). Nevertheless, this use is not what fixes the utterance's meaning, according to the invariantist. The act of understanding the utterance, for the invariantist, is a simple matter of taking it to mean something like (5). The hearer(s) then situate(s) the (implicitly or explicitly recognized as false) utterance in the appropriate context so that it comes out as a worthwhile thing to say under the given circumstances. The invariantist, in other words, will make sense of the utterance in a way that answers the question of why the speaker said something that is, strictly speaking, false, by giving the falsehood a *non-semantically relevant role* of, e.g., fixing everyone's gaze on a place that would make for a great picnic. The contextualist, on the other hand, will pack the requirements for serving such purposes into the very meaning of the utterance, rendering it strictly speaking *true*. In sum, the difference is that an invariantist will claim that people often utter falsehoods for good pragmatic reasons, whereas a contextualist will claim that since those same pragmatic reasons fix meaning, what are uttered in such cases are oftentimes truths.

So, to return to the relevant matter of how to construe claims to know, when one makes such a claim, according to the contextualist, one is stating that one is in an epistemic position vis-à-vis the claim one purports to know that is strong enough to meet

the standards for certainty currently in place. It is important to see, then, that the contextualist does not claim that the truth of a knowledge claim depends *only* upon the strength of one's epistemic position. Depending upon the standards in place, one and the same type of epistemic position may or may not result in a true knowledge claim. Just as how in (3) what the referent is of the singular term "that field" can remain constant even though the truth value of (3) fluctuates with the relevant change in context, so too can the referent of "so-and-so's epistemic position" remain constant even though the truth value of "So-and-so's epistemic position constitutes knowledge for the claim that p" fluctuates in like manner. For the contextualist, then, it is not *just* one's epistemic position that determines the truth value of a knowledge claim, but also how good one's epistemic situation must be--i.e., how *certain* one must be--in order for such claims to count as knowledge that plays a determining factor.

Now, as it just so happens, the contexts of everyday assertions to knowledge are usually lax enough so that assertions like (1) come out true. After all, to be certain that Barack Obama is currently President requires, for everyday contexts, nothing more than that one be apprised of current events in some socially accepted and trusted (perhaps because of its track record) fashion. In this way, the contextualist thinks they can reconcile the fact that we know things like (1) in most normal contexts with the persuasiveness of what seems to follow about our knowledge granting fallibilism. The idea is that *as soon as the skeptic conversationally brings up the skeptical scenarios*, the standards for certainty are raised accordingly. In these cases the contextualist (as a good fallibilist) claims that the standards become so high that we fail to know claims like what

is found in (1). But since there are other, more everyday contexts where the standards are lower, we also oftentimes know plenty about the external world (DeRose 1995, 185-187).

The contextualist's position presents, as I see it, an extremely odd iteration. One would think that the question of whether or not one possessed knowledge of a particular fact, for example, depended only upon what one's epistemic position *in fact was*; but the contextualist, as we have seen, claims that one's epistemic situation can remain the same, while one's claims to know become either true or false as a result of the shifting of (purportedly) conversationally-fixed standards for certainty. But why aren't the latter factors simply irrelevant? It would seem that if I was in fact in a strong epistemic position with respect to some claim, then the last thing that would affect my ability to legitimately call my epistemic position "strong" would be something someone else (or myself when I adopted a skeptical mood) *said*. What if the person saying something (including the skeptical version of myself) is a complete fool? Obviously, we shouldn't take them seriously. But this means that what we, just as a matter of fact, take seriously is a function of how we *judge* what is in fact being said. And it seems, though I will have to make this clearer, that our basis for so judging would be nothing other than *what we take our own epistemic position to be* with respect to the claim in question.

This last point provides the source of my main criticism of the contextualist's position: it rests on a confusion about what really motivates the problem of philosophical skepticism. First let me say some things that merely state the pessimistic fallibilist's reaction to the contextualist view; arguments will follow. According to the methodology that is the sort of first-person, neutral inquiry sketched above, the skeptic does not hold court because he seduces us into implicitly and contingently accepting his high standards.

He holds court because he questions the truth of our contingently held belief--made evident by our willingness to make the relevant assertions--that our epistemic situation vis-à-vis external world claims is in fact favorable to at least *some* degree. The problem of philosophical skepticism arises because if we are fallibilists who adhere to a neutral, first-person methodology of questioning such matters, this contingently held belief can be proven to be false. This is done by proving how fallibilism entails that the (likely) truth of external world claims is, in the sense I have discussed repeatedly, presumptuous at best. In this way the skeptic is *most accurately* understood as questioning the truth of our beliefs regarding the referent of the singular term "my epistemic situation" in a claim like (2), the very referent that, according to the contextualist, can remain constant even as (in their view) the truth value of the relevant knowledge claim shifts.

So, to make the confusion about the source of the skeptical problem evident, the pessimistic fallibilist should simply ask the contextualist what he/she takes to be in certain cases the (ex hypothesi) stable referent of the singular term "my epistemic position." As we have seen, the only ones who will think there is a problem of philosophical skepticism as I understand it will be those who take the referent in question to be a *favorable* one regarding certain external world claims. Obviously the contextualist will have to agree that we maintain this attitude even in the most lenient of contexts. After all, since "know" is a positive term of appraisal, it follows that even in lax contexts one's epistemic position must be *somewhat* favorable in order for the claim to come out true. But, as we have seen, if you are a fallibilist then for familiar reasons you must admit that it follows that your epistemic position is not favorable at all. So, if you are a *pessimistic* fallibilist, then you have to feel the force of the problem of philosophical skepticism as I have

framed it. But, I have just argued, if you are a contextualist you have to do so as well. This is because it is clear that our positive attitude towards the status of the referent of "my epistemic position" with respect to external world claims is necessarily a matter, if not more fundamental, then at the very least inseparable from what allows (says the contextualist anyway) the knowledge claims to be true in even the most lenient of contexts. But what is equally clear is that it is *merely* the interaction between what we take our epistemic position to be and what follows from fallibilism that generates the problem of philosophical skepticism. How conversationally-determined standards for knowledge might shift in this or that way is a matter that is therefore completely orthogonal.

For this reason, the contextualist cannot make the claim that they often take to be the greatest virtue of their position: that their view provides a way to both explain the force of the skeptic's reasoning, while also preserving the plausibility of common sense epistemological realism. If I am right, then at the very least the contextualist fails horribly at the former task. And, of course, if the pessimistic fallibilist is right, then in a qualified way they fail at the latter task as well.

Sweeping remarks like the above will not convince one who is familiar with particular, well-argued for contextualist positions, and has thereby become convinced of their intuitive force. So, in order to drive the matter home I will now critically examine a prominent contextualist view; namely, the one espoused by Keith DeRose in his (1995; all of the page references that immediately follow refer to this work).

DeRose thinks that a central concern of the contextualist should be to explain what the actual mechanism is that drives a shift in standards for certainty. In other words,

DeRose thinks, nobly enough, that a contextualist is called upon to explain how it is that the standards might be raised, lowered, etc., in a particular conversational context. His view about how the shifts often work stems from what he takes to be a plausible explanation of our intuitions about the status of various knowledge claims. As he writes:

...we have a very strong general, though not exceptionless, inclination to think that we don't know that P when we think that our belief that P is a belief we would hold even if P were false. Let's say that S's belief that P is *insensitive* if S would believe that P if P were false... We tend to judge that S doesn't know that P when we think S's belief that P is insensitive (p. 193).

It is important to see that, generally speaking, what beliefs are insensitive in this way is a function of the strength of one's epistemic position. It seems correct, for example, that most of us do not *know* that the zebra-looking animals we find in zoos are not cleverly painted mules. DeRose's favored explanation for this failure to know is that in the context of looking at the relevant animal in broad daylight while, say, it mills around in its pen, our belief that it is not a cleverly painted mule is insensitive. Most of us, in other words, would still believe that any zebra-looking animal we see under such circumstances was not a cleverly painted mule, even if it was in fact a cleverly painted mule.

Nevertheless, matters are quite different for a zoologist. A zoologist presumably knows enough about how both zebras and mules look to tell by sight alone that any zebra-looking animal is not a cleverly painted mule, and that a mule cleverly painted to look like a zebra is in fact just a mule. Unlike most folks, if a zoologist was presented with a mule that was cleverly painted to look like a zebra, he/she would *not* believe that he was not looking at a cleverly painted mule. The zoologist's belief here is, unlike most of us, not insensitive, and this is because the zoologist is in a stronger epistemic position than us with respect to the belief that the animal in the pen is not a cleverly painted mule.

DeRose uses this sensitivity intuition to explain how, he thinks, the skeptical problem arises. He maintains that we intuitively believe that we know, for example, that we have hands because we take ourselves to be sensitive to that claim's falsehood (p. 197). But equally true is that we believe that we do not know that we are not dreaming, not a brain in a vat, etc., because we take ourselves to be insensitive to the falsehood of these claims (p. 194). These two beliefs about what we know/do not know, combined with the extremely plausible claim that knowledge is closed under known entailment, produce a paradox (one different, albeit related, to the paradox that arises according to the pessimistic fallibilist). According to DeRose, as I've just explained, each of the claims

(7) I do not know that I am not dreaming.

(8) I do not know that I am not dreaming only if I do not know that I have hands.

(9) I know that I have hands.

enjoys initial intuitive plausibility (p. 183). And yet, of course, they all can't be true. Our making use of the inherent plausibility of a requirement for sensitivity is DeRose's preferred way of stating what causes us to come to accept (7) and (9), thus aiding in creating the problem of skepticism.

What motivates (8)? To explain this aspect of DeRose's view, I must explain what he takes to be a workable enough method for determining *relative strength of epistemic position* with respect to certain claims. In certain cases, DeRose claims, we can take the truth of a conditional like *S knows that p only if S knows that q* to depend upon what we recognize to be facts about S's comparative epistemic position vis-à-vis the claims p and q found within that conditional. So, DeRose claims, we recognize the truth of the conditionals (8) and its contrapositive:

(10) I know that I have hands only if I know that I am not dreaming

because we recognize that we are in at least as strong an epistemic position with respect to the claim that we are not dreaming as we are in with respect to the claim that we have hands. As DeRose writes:

Given natural background assumptions, we can sense that the following comparative fact holds...I am in no better position to know that [I have hands] than I am in to know that [It is not the case that I am dreaming]. This comparative fact is revealed...by the highly plausible conditional that...: If I don't know that [I am not dreaming], then I don't know that [I have hands]. Closely tied to that comparative fact...is the related and intuitively compelling realization that it would be no wiser to bet one's immortal soul on [the claim that I have hands] being true than to bet it on [the claim that I am not dreaming] being true (p. 203).⁴²

Given this, plus the previous results, we are in a good position to notice that "One's epistemic position with respect to propositions to the effect that skeptical hypotheses don't hold must be stronger than it is with respect to other, more ordinary propositions if belief in such propositions is to be sensitive" (p. 204). This falls out of the dual realization that sensitivity is a function of strength of epistemic position, and that my belief that I'm not dreaming is insensitive. DeRose thinks we all realize that we are not sensitive to the falsehood of the skeptical scenarios, and also that we *are* sensitive to the falsehood of beliefs like that we have hands. So, since sensitivity is a function of strength of epistemic position, it follows that in order to make the belief that we are not in a skeptical scenario such that we *are* sensitive to its falsehood, we must get ourselves to be in a stronger epistemic position than we are ever required to be in order to know that we have hands.

This last point, DeRose claims, "suggests a new contextualist account of how, in presenting [the skeptical scenarios], the skeptic raises the standards for knowledge" (p.

205). Here finally, we get at how DeRose thinks changes in standards for certainty can take place in certain conversational contexts. In general, we all obey what DeRose calls a

Rule of Sensitivity:

When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards [for] how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular P to be sensitive for it to count as knowledge. Where the P involved is to the effect that a skeptical hypothesis does not obtain, then this rule dictates that the standards will be raised to a quite high level, for, as we've seen, one must be in a stronger epistemic position with respect to a proposition stating that a skeptical hypothesis is false--relative to other, more ordinary, propositions--before a belief in such a proposition can be sensitive (ibid.).

Thus, according to DeRose, merely when the skeptic asserts that we don't know that we are not dreaming, he raises the standards for knowledge

to such a level as to require our belief that [we are not dreaming] to be sensitive before it can count as knowledge. Since our belief that [we are not dreaming] isn't sensitive, the standards are driven up to such a level that we don't count as knowing that [we are not dreaming]. And since we are in no stronger epistemic position with respect to [the claim that we have hands] than we're in with respect to [the claim that we are not dreaming], then, at the high standards put in place for the skeptic's assertion...we also fail to know that [we have hands]. At these high standards, the skeptic truthfully asserts...that we don't know that [we have hands]. This accounts for the persuasiveness of [the skeptic's claim]. But since, on this account, the skeptic gets to truthfully state her conclusion only by raising the standards for knowledge, [this conclusion] doesn't threaten the truth of our ordinary claims to know [things like that we have hands]...For the fact that the skeptic can install very high standards that we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates (pp. 206-207).

I hope this gives the reader a workable feel for DeRose's response to skepticism.

Unfortunately, it predictably falls prey to the general problem I have raised for the contextualist. DeRose seems correct in pointing out that intuitively we do not think that we know something if we think we are insensitive to the falsehood of the relevant claim. But he fails to realize the implications of the fact that which claims we think we are

sensitive to is a function of what we take to be the strength of our own epistemic position. The point about the difference between the zoologist and the layman with respect to the claim that a zebra-looking animal is not a cleverly painted mule, shows quite clearly how we cannot answer the question of which beliefs we are, *in fact*, sensitive to without *first* (or at least also) answering the question of how strong our epistemic position in fact is. It follows that according to DeRose's own view of matters, in order to understand how the mechanism of the Rule of Sensitivity might operate in any context, we must first (or at least also) determine how strong the subject in question's epistemic position *in fact* is. Only once we have answered this question can we determine if shifts of standards for knowledge do indeed ever operate as DeRose claims they do when the skeptic comes on the scene. But, as I've already pointed out, the pessimistic fallibilist has a way of motivating the problem of philosophical skepticism that *only* depends upon our taking ourselves to be in some sort of favorable epistemic position with respect to certain external world claims. If we are also fallibilists (and DeRose certainly is), then we must *also* come to think we fail to be in a favorable epistemic position for such claims. Thus, it is clear that we can explain the *source* of the problem of philosophical skepticism in a way that has nothing to do with the matter of how conversational context might shift in this way or that. And, moreover, we have seen that according to DeRose's own position regarding how Rule of Sensitivity governed shifts of standards work, whether things operate in a way favorable to his resolution of the skeptical paradox will depend on what our own epistemic position in fact is. But, the important point is, as long as we *merely take it* to be positive in *some* sense with respect to external world claims, the pessimistic

fallibilist can make his position regarding the problem of philosophical skepticism abundantly clear. How *shifts* in standards work is completely orthogonal.

This problem with DeRose's account can be brought out in a different way. It is of course a truism that we all *take ourselves* to know things like that we have hands. But what I fail to understand is why DeRose is so confident that we also all take ourselves to not know that we are not dreaming. I think it is true enough that most folks think a claim like (10) holds, and I am willing to concede that this reflects the belief that we are in at least as strong an epistemic position with respect to the claim that we are not dreaming as we are in with respect to the claim that we have hands. But this is also consistent, of course, with thinking that the converse of (10) holds as well; i.e., that we are also in at least as strong an epistemic position with respect to the claim that we have hands as we are in with respect to the claim that we are not dreaming. In other words, maintaining (10) is consistent with thinking that the strength of our epistemic position is *identical* for the two claims. Perhaps knowing that I am not dreaming is a matter that stands or *falls* depending upon whether or not I know *things like* that I have hands. Certainly knowledge claims of this "Moorean" type have a special status.⁴³ I may, in other words, know that I am not dreaming while forever failing to know, say, facts of history or geography. And the reason may be that I still, nonetheless, know only things like that I have hands, allowing me to infer that I am not dreaming. In sum, perhaps it is correct that:

(11) I know that I have hands *if and only if* I know that I am not dreaming.

Perhaps, that is, the *only kind* of knowledge that would allow me to know that I am not dreaming is my knowing things like that I have hands. In any case, the important point is just this. Anyone who maintains (11) would, contra DeRose, deny that we must be in a

stronger epistemic position in order to know that we are not dreaming than we must be in order to know that we have hands. Thus, on this view, by bringing up the skeptical scenarios the skeptic does *not* raise the standards in the way that DeRose claims. Of course, the same person may or may not think that they were sensitive to the falsehood of the claim that they have hands. But, no matter what, they would have the same view for both claims here. This shows quite clearly how something inseparable from the task of determining how shifts in the standards of certainty actually work, is a determination of *what grounds* we have for thinking that our epistemic position is of this or that strength with respect to some set of claims. We *cannot* say that things shift in the way DeRose wants to say they do without *first* (also) determining the strength of our epistemic position.

This shows how DeRose's view is orthogonal in a way that I suggested all contextualist positions will be. The problem of skepticism arises from an inconsistency between what we take our epistemic position to be (something made evident by how, as DeRose would put it, we take ourselves to know things like that we have hands) and what fallibilism entails it to be. If fallibilism is correct, then our epistemic position is in fact weak in the sense I have pointed out again and again. If we believe in fallibilism, then we must believe in this kind of weakness. But this is inconsistent with our equally strong belief that our epistemic position is strong (at least to some degree) for certain external world claims. DeRose's explanation of the problem of skepticism does not work, then, because it is not at first sight clear that we all need to accept, to give just one example, (10) but not (11). Our views on this matter, generally speaking, will depend upon what we take our epistemic position to be. Only by way of a certain understanding of our

epistemic position will the heart of DeRose's view gain traction. The upshot is that even if we take DeRose to be on the right track the important question *still* becomes: How strong is our epistemic position? But, as my arguments have hopefully made apparent, when we realize simply that we *take ourselves* to be in a *somewhat* strong epistemic position with respect to claims about the external world, and also that fallibilism sounds intuitive, the problem of philosophical skepticism presents itself in full.

3.4 Process-Reliabilism

I will now move to a critical discussion of the response to skepticism that might be given by a process-reliabilist. I should emphasize the term "might" here. Historically speaking, process-reliabilism began being widely discussed as a result of Alvin Goldman's desire to deal with Gettier counter-examples to a popular analysis of knowledge which understands it to be *justified true belief* (see Gettier 1963; Goldman 1967). Goldman has also, at times, sounded indifferent with respect to the question of whether or not a reliabilist can adequately deal with skepticism (Goldman 1986, 40-41). Thus, it should be made abundantly clear at the outset that no process-reliabilist need be *bothered* by the fact that their view is compatible with having to admit that the problem of philosophical skepticism exists as I have construed it. Nevertheless, seeing why a process-reliabilist can't answer or deal with the skeptic will be instructive for various reasons, not least of which is how it helps to further indicate the power of the problem of philosophical skepticism. Also, as will be shown in due time, their view seems to point to features of our external world beliefs that one who wishes to refute the fallibilist must ably illustrate to be present. So, let me now explain the basics of the reliabilist position and then get into how a reliabilist consequently would react to the pessimistic fallibilist's

view. (I should note that the discussion that follows mostly pertains to the early work of Alvin Goldman, specifically certain claims found within Goldman (1967, 1976)).

Intuitively, a reliabilist thinks, a justified belief is one that has a high *objective* probability (in some sense of that word) of being true. For the process-reliabilist, more accurately, any belief is justified when it is the result of a reliable belief formation process. Such belief formation processes are of two main types: *belief-dependent processes* and *belief-independent processes*. Belief dependent processes are those which produce beliefs via a process that has a subject's prior held beliefs as input. The obvious example is inference. Sound deductive inference, for example, is a paradigmatically reliable belief-dependent belief-formation process. Deduced beliefs, generally speaking, are those which depend, for their formation, on the presence of prior beliefs that serve as the premises for the deduction, and thereby the inputs in the process which produces the deduced belief.

Belief-independent belief formation processes, on the other hand, are those where the process involved does not depend, for its operation, on the presence of a prior belief that the subject holds as an explicit input. The paradigm of such a process, for the process-reliabilist anyway, is perception. For the belief-formation process of perception, what is required as an initial input is the non-propositional perceptual experience that is, say, a particular visual or auditory experience. I form the belief that a computer is in front of me, it seems, primarily as a result of the fact I am having the relevant visual/tactile/auditory experiences. No beliefs play a straightforwardly *obvious* role in what causes the formation of this belief. According to a *crude* reliabilist position, then,

just as long as a belief-formation process is in fact suitably reliable, the output beliefs it generates are consequently epistemically justified.

Seeing where the crude view leads will be helpful. We can now ask, that is: how can a view which understands justification *solely* in terms of reliable belief formation processes respond to the pessimistic fallibilist's position? It can do so by simply admitting, first of all, that most people will be unable to refute the claim that they have no evidence for their external world beliefs, and then pointing out that this alone has no bearing whatsoever on the justificatory status of our external world beliefs. My admitted inability to argue effectively against someone (including myself when I reflectively ask questions about the epistemological status of my own beliefs) in and of itself, has no bearing, says the crude sort of reliabilist, on whether or not my perceptually-based beliefs are justified (because reliably-based). The reliabilist response, then, is that an *intuitive understanding* of the concept of justification ties it to reliability, and hence divorces it entirely from the requirement that we be able to refute a certain person holding a negative epistemological claim. So, even if our beliefs are presumptuous at best in the sense I have maintained, as long as they are in fact reliably-based they are still nonetheless justified. The arguments of the preceding chapters, then, in of themselves cannot establish anything conclusive about our negative epistemic position with respect to external world beliefs, and for that reason they are at best not worth troubling over.

My response to crude reliabilism on the behalf of the pessimistic fallibilist will involve going deeper than just repeating the point that the pessimistic fallibilist cares nothing about whether or not they wield an intuitive understanding of justification, and is quite legitimately skeptical of the value in performing armchair conceptual analyses.

Though, to be sure, I think that these points alone shows why the reliabilist response to skepticism is in general without merit, I also think there is a much more troublesome defect in reliabilism, crude or sophisticated, in this regard. It is particularly evident for the crude form of reliabilism, and thus it will be helpful to begin by explaining how it arises for that position. It has to do with what falls directly out of another point I have been emphasizing ad nauseum: how when we make assertions we contextually imply that we find ourselves adequately evidentially backed. *All* reliabilists are, I take it, willing to make various assertions about the external world. But, if so, then all reliabilists more or less *explicitly/directly* contradict themselves; and what is so interesting is to witness how they cope with this situation by adjusting matters accordingly, as they do, we shall see, in the literature. First, however, let it be made clear that crude reliabilists in particular will, in making external world assertions, continuously perform an act which in effect implies that they think they *can* effectively answer a higher-order skeptical worry (which according to their view would require them to *show*, in particular, that their belief is reliably-based), but then openly admit, and even *embrace* the fact that they *can't!* (There is, however, an important caveat to mention here; one that I will address in just a moment.) Importantly, because they not only admit but even *embrace* the fact that they can't answer a higher-order worry, and yet are also willing to make the relevant assertions that thereby contextually imply that they *can* answer a higher-order worry by *showing* that a belief is (likely to be) true, they maintain a view that should allow them to easily notice what I have called the problem of philosophical skepticism. And yet, not only do they not seem to think there is a problem of philosophical skepticism, because they *embrace* rather than, say, merely grudgingly admit their inability to effectively answer a

higher-order worry, they reveal an *extremely* peculiar feature of their position. They respond to the paradox by, in net effect, embracing the *naturalness* (!!) of an utterance like: *There is soda in the fridge and I am in no position to ably counter someone who thinks that there isn't soda in the fridge*. If the crude reliabilist is like most of us, he will be willing to make the first assertion under the appropriate circumstances, thereby contextually implying that he takes himself to be in a position that directly contradicts the statement that follows the "and"--even though that same statement is also (in a qualified way) something that he *embraces* given his own philosophical position vis-à-vis higher-order requirements. Something has gone terribly wrong here.

Let me now deal with two objections that immediately arise. First, it is not *quite* correct to say that a crude reliabilist believes that they *can't* answer a higher-order skeptical worry. An effective answer to a higher order skeptical worry, according to crude reliabilism, would be an appropriately formed meta-epistemological belief that was also, as it just so happens, reliably-based. But, of course, the mere confident statement of a meta-epistemological belief, reliably-based or otherwise, would not work to meet the argumentative challenge of one who doubted that your original, first-order belief was reliably-based. The reliabilist, as a good fallibilist, does not think that we could quiet such a challenger by showing them that our beliefs are, in fact, reliably-based to even some appreciable degree. At best they maintain that they can say to their challenger that *hypothetically* their external world beliefs are reliably-based. This is the sense in which a crude reliabilist must admit that they can't effectively answer a higher-order skeptical worry. And, for my purposes, it is the important sense. After all, when any crude reliabilist makes an assertion about the external world, he contextually implies that he

can do just that--where in his case this means he implies that he can show that his belief is in fact reliably-based. In performing that act, then, he thereby contradicts himself in the relevant sense.

The point is worth emphasizing in a different way: it is absolutely crucial to see that, when we actually make assertions, we take ourselves to have evidence in a sense *identifiable with* thinking that (given its basis) a belief is *in fact* (likely to be) true. It is completely incorrect to maintain instead that, for example, when we make assertions we take ourselves to have evidence in the sense that we think that (given its basis) a belief is--just as long as we also happen to be right about the reliability of its basis--(likely to be) true. Rather, we take ourselves, in such cases, to have an *effective* answer to the question: "What makes you think that's (likely to be) true?" It is an answer that we believe will *effectively* show why someone who doubted us is (likely) wrong. We do not take ourselves to have a *hypothetically* effective answer in this regard. Importantly, this is confirmed, quite conclusively as far as I can tell, by the fact that if after making an assertion we find that we *can't* come up with an effective answer of this sort, we either retract the original claim or else subsequently adopt a weaker attitude towards it. This shows well enough, I think, that we (implicitly) took it to be a precondition for making the original claim that we be able to come up with an effective answer. Obviously, if we instead took it to be a precondition for making an assertion that we just have a *hypothetically* effective answer, then we would not weaken or retract the claim in question upon realizing that it is perfectly within the bounds of intelligibility that, at that time, we are not effectively answering it. This is not what happens. It is due to the simple fact that when we make assertions we do *not* take ourselves to have a *hypothetically*

effective answer to some challenger, then, that the sense in which the crude reliabilist *can* admittedly answer a higher-order worry is irrelevant.

The next objection brings to the fore the issue that the above development stirs up. It might be claimed that my account of what goes on when we actually assert things about the external world confuses the phenomenon of (by contextual implication) taking oneself to have good evidence in the strong sense I require for what is in fact just one's taking oneself to, as it is often put, *for-all-practical-purposes* have good evidence for a claim. This would attempt to remove the above absurdity by saying that what a crude reliabilist (and others) is (are) *really* willing to claim is something more like: "p (or, at least, this is what I am sensibly maintaining as true in order to get on with the business of everyday life) *even though* of course I am not positioned to ably counter the claim that not p." Thinking we contextually imply that we are justified for-all-practical-purposes in this way when making assertions is, of course, quite different from thinking that we contextually imply that we are positioned to *actually* refute someone who doubted us. In order simply to get on with normal life, according to this view, we go about our day to day affairs by, in effect, *assuming* that our epistemic position is strong with respect to certain claims--thereby often claiming, and even sometimes insisting, that various things are, in fact, (likely) true. We assume so because it is the sensible thing to do given that we want to act effectively, thrown into the hustle and bustle of daily life as we are. People just don't have time to first carefully think through the amount of evidence that they have for this or that claim about the external world, and *then* act accordingly. We just have to act, and this includes making assertions. Nevertheless, it is also true that upon proper reflection we can realize that, for *anything* we say about the external world, as it

turns out we *cannot* effectively argue against various kinds of challengers who deny the (likelihood of the) relevant claims. But there is nothing contradictory in maintaining this dual attitude. Given the central role that I have our willingness to make assertions about the external world play in my understanding of what generates the skeptical paradox--in that I have been insisting that most of us are, as a matter of fact, willing to make such assertions in the strong sense I require--it is obvious that getting clear on this issue is of paramount importance.

Attempting to draw a distinction between what we do in day-to-day life versus what we do as philosophers is an old tactic for coping with the skeptical paradox. It goes back at least as far as the writings of Hume, who famously said things like "Be a philosopher; but, amidst all your philosophy, be still a man" and "As an agent, I'm quite satisfied in the point; but as a philosopher...I want to learn the foundation of this inference" (Hume 1748, pp. 90, 117). For my own part, however, I must say that this distinction has always struck me as spurious. It is true enough that in order for us to get on with the business of everyday life we must, at times, *take ourselves* to be in a strong epistemic position with respect to certain external world claims. But isn't it just as much a part of so taking ourselves to maintain that, *whatever* might show our epistemic position to be, in fact, weak--forcing us to revise what we take our epistemic position to be--*the last thing* that would do so is the mere words of any old challenger? The point goes back to a criticism of the contextualist that I made earlier. We *never* allow what we take our epistemic position to be to be affected by the mere protestations of someone we *deem* to be a complete fool. Instead, generally speaking, which sorts of considerations we allow to affect the matter of what we take our epistemic position to be is a function of what we

judge to be the *merit* of the contrarian claims offered. *This is just as true in everyday life as it is in the most abstruse of philosophical contexts.* What lies constant throughout is what sorts of things we, in fact, allow to affect what we take our own epistemic position to be with respect to certain claims. We should now ask, then, what is the *merit* of the following suggestion: *what we take our general epistemic position to be with respect to the external world is something that should oscillate with context?* And here I can begin only by merely stating my own view. It is utterly *mystifying* to me why anyone would think that even though a consideration offered in a philosophical context was deemed plausible enough at that time to affect what one took to be one's general epistemic position with respect to external world claims, it somehow loses any of its ability to do so in a more everyday context, when we are not considering such questions, and where immediate action is required. What sense is there in allowing the mere fact that someone is no longer talking to you (or that you are no longer talking to yourself) to affect what you take to be your general epistemic position with respect to external world beliefs? What, in other words, is the *merit* of the following point? Before, when you were in the marketplace it was perfectly okay for you to take your general, overall epistemic position with respect to external world beliefs to be strong, but now that I'm talking to you and have said various things it is no longer okay for you to do the same. Is this point of any merit, on the face of it? Even the contextualist, as far as I can tell, agrees that what our epistemic position in fact is, is not something that changes with conversational context. Therefore, equally unaffected, I would think, *should* be what we take our epistemic position to be when we oscillate between various contexts. Whatever you deem to be your epistemic position as a philosopher should, for this reason, be the same as what you

deem to be your epistemic position "as an agent," and vice versa. If nothing else, this should be because you should see how what changes as we oscillate between these contexts is completely *irrelevant* to what your epistemic position with respect to external world claims, generally speaking, in fact is.

At this point I imagine the following protest: "Well, if we didn't simply assume that we knew (or had reasonable enough grounds for maintaining) plenty about the external world in everyday life, wouldn't this stifle practical activity? Why, for example, would I go to the store to buy groceries if I also thought that I had no good reason for thinking that the store even existed? Generally speaking, moreover, if we don't think we have good reason for believing that something in the future will be true, we do not act presently as if it will be true. For this reason, isn't it *best* that, at certain times at least, we ignore our abstruse, philosophical conclusions and just get on with life?"

But the very statement of this objection, to my ears, only serves to bring to light the *paradox* at hand. It points squarely to why this objector cannot themselves *really* accept the reasonableness of the claim that our epistemic position is in fact weak with respect to external world claims--the very thing that, however, the truth of fallibilism entails when we undergo a first person, neutral inquiry of the extent of our knowledge. The conflict between what we want to say "as an agent" and what want to say "as a philosopher" only illustrates that there is the paradoxical mismatch here regarding what we think; and what my latest objector only confirms is that the paradoxical nature of this mismatch is not going away. Thinking "as an agent" that, for example, I cannot show someone why they are wrong by allowing them to *see for themselves* is certainly the

height of absurdity (right?). And yet, to accept fallibilism "as a philosopher" is to treat the absurdity as an obvious truth!

Aside from the above more or less rhetorical remarks, I can thankfully respond more substantively to the last objection made by combining some previous remarks with the epistemological insights that can be drawn from Wittgenstein. This will provide the theoretical considerations I wish to bring to bear to support the above mere declarations. The distinction my objector wishes to draw here runs afoul of the Wittgensteinian point that was summarized above as the claim that *epistemological beliefs presuppose non-epistemological beliefs*. The phenomenon of *merely assuming* something substantive regarding our epistemic position when we make an assertion while involved in day-to-day affairs is best seen as, instead, simply behaving in accordance with an adopted logic of inquiry, which results in certain things *at a particular time of action* being (perhaps just temporarily) beyond doubt. What my objector treats as taking ourselves to have good evidence for-all-practical-purposes is, rather, us taking ourselves to have good evidence *simpliciter* according to a common sense world-picture. While adopting that world picture, we just, in point of fact, (implicitly) treat our epistemic position with respect to certain external world claims as *strong*--not as hypothetically strong, or strong for all practical purposes--to be such that if there was a doubter of the matter we could *show him* why he is incorrect by undergoing a certain procedure. And, as Wittgenstein ably pointed out, inextricably involved in so treating matters is the treating of certain claims as beyond doubt. In sum, for Wittgensteinian reasons we can say that my objector mistakes our willingness to assert something outright according to a certain logic of inquiry for our willingness to assert something as true for-all-practical-purposes. So, if a reliabilist is

willing to assert something like that there is soda in the fridge in a normal context, then this does imply, contra what my objector maintains, that he takes himself at that time to be able to *show* a challenger why they are incorrect. When we combine this point with what I've already declared, which is that our attitude with respect to our own epistemic position *should* not change when we are doing philosophy versus when we are involved in everyday life, this illustrates quite clearly why the crude reliabilist, of all people, should be the first to notice the existence of the skeptical paradox. But, the upshot of my response to the latest objection is confirmation of the fact that the reliabilist does seem required, rather, to embrace as natural, the sorts of claims that make the paradox evident.

This last objection does, however, raise the legitimate (and by now easily anticipated) worry that I am hanging too much on the contingent fact that we are willing to assert things about the external world, in a sense that implies that we take ourselves to be able to *show* that the claims in question are (likely to be) true. Given the amount of work that this claim is doing for me, it seems I am called upon to positively argue for it much more than I have up to this point (which is not much at all). At the moment, however, I can only promise that an argument for the much *stronger* claim that it is *reasonable* to assert things about the external world in the relevant sense will eventually be given--specifically in the next chapter. Until then, however, it must be admitted that it is, to the extent that such an argument is still deemed to be required, just as much an open question that the crude reliabilist is *forced by rational considerations* to be compelled to assert things outright about the external world in the strong sense I require--and thus forced to embrace the contradictory nature of their overall position. For now, however, I

will continue to treat them, not as being so forced, but rather just as a matter of *contingent fact* being willing to do so.

I have admitted, in any case, that the position I have been operating with is a crude reliabilist one. A more sophisticated form will, of course, make at least some concessions to higher-order epistemological requirements. Specifically, a more sophisticated form of reliabilism will admit that because we are, by default, willing to make assertions in the sense of treating ourselves as able to show why some challenger is wrong, it follows that we would never actually assert something like: "p and I'm unable to ably counter someone who maintains that not p." A more sophisticated form of reliabilism, in other words, would not embrace the *naturalness* of such assertions--though I hope I have shown that this, oddly enough, is what a crude form of the view is forced to do. Instead, a more sophisticated form of the view will claim that when we are presented in a certain context with skeptical reasoning our own *subjective rationality* requires us to stop being willing to make assertions about the external world. But the sophisticated reliabilist will wish to emphasize that this is compatible with us being completely justified in maintaining that p when we are in fact able to reliably assert it (i.e., in cases where factors having to do with subjective rationality do not block our ability to do so) *and* no other factors that could just as well get in the way obtain--for example, the possession of reliably-based countervailing beliefs.

What is still difficult to understand, however, is how this kind of reliabilist, being a fallibilist, will not be required to agree that the essence of the problem of philosophical skepticism as I have framed it exists in full force. Take, for example, a certain series of stages in the storied evolution of Alvin Goldman's defense of reliabilism. In an early and

influential paper, Goldman states the predictable reliabilist claim that tries to separate being "really" justified from (in effect) being able to refute a skeptic. Drawing on an oft-deployed analogy with ethics, Goldman makes the following distinction:

On the one hand, a principal of justification might specify the features of beliefs (or other doxastic attitudes) that confer epistemic status. These features may or may not be usable by a cognizer to make a doxastic choice. On the other hand, a principal of justification might be designed specifically to guide a cognizer in regulating or choosing his doxastic attitudes. Here the criteria of justification must be ones to which a cognizer can appeal in the process of making a doxastic decision. That the theoretical and regulative functions of justification principles can be distinct emerges clearly from an account of justified belief I propose in another paper [Goldman 1979]. Refinements aside, this account...says that a belief is justified just in case its causal ancestry consists of reliable belief-forming processes, i.e., processes that generally lead to the truth. As a theoretical specification of epistemic status, such an account is entirely suitable. But this theory or principal cannot be used by a cognizer to make a doxastic decision; nor is it so intended (Goldman 1980, 37-38).

Here it is important to note that what would in *every case* serve to guide a cognizer with respect to whether or not they themselves *should* hold that a certain belief *is (likely) true* is nothing other than whether or not they are able to effectively answer skeptically-focused questions of various sorts, and of various levels of extremity and depth. It is the broadness of scope in what it is to be a "skeptically-focused question of a certain level of extremity and depth" which allows this last point, at least upon examination, to amount to a simple truism. Therefore, by recognizing that a cognizer can seek principles which are supposed to guide his *cognition*, Goldman in effect recognizes just that there is an enterprise for skeptical questioning (very generally understood) in epistemology. Combining this realization with Goldman's point, we get the following variation-by-terminological-gerrymandering of something I've stressed earlier: that when you proceed with a first-person, neutral methodology, what it means to be "really" justified is just identifiable with which beliefs survive skeptical attack. Given Goldman's desired

distinction between being "really" justified and being what can actually guide a cognizer when forming their beliefs, that is, we must merely re-state the claim as follows: when we are concerned with determining which of our own beliefs we should/should not think are true, we must proceed by determining which among them survive various kinds of skeptical attack.

But, as the points I've made against crude reliabilism effectively indicate, it is clear that you cannot keep matters so cleanly separated--between what Goldman calls "theoretical" versus "regulative" principles of justification. Clearly enough, for example, no cognizer would *believe* that p even when p was a claim that was, in fact, reliably-based (say, by perception) when they also, *for various reasons they reflectively deem to be quite convincing*, believed (falsely) that p was either false or highly objectively unlikely. You can't believe something unless you think it is true. So, obviously, if no one would ever believe that p under such circumstances, they certainly could not *justifiably* believe it. The fact that it was reliably produced by perception would in this case be idle, epistemically-speaking. This indicates well enough that there simply *cannot be* claims one justifiably believes that are also claims that are *subjectively irrational*--and this is so, even when the claim in question is in fact supported by a certain reliably-based process. Now, to at this point try to cling to the essence of one's original reliabilist position--by here emphasizing the distinction between being propositionally justified and doxastically justified, claiming that reliability does still pertain to the former property of a belief even when that belief is subjectively irrational--is, for *my* purposes, at best a needless verbal filigree. The fact remains that there *cannot be* subjectively irrational claims one justifiably believes; again, simply because one would never believe such claims at all.

Furthermore, and much more importantly given my own purposes, for the very same reasons no one would *assert that p* under such circumstances, even when it was in fact reliably-based by, say, perception. The latter fact would, in such cases, serve no epistemic good. This shows that what *sometimes* has to be at least *in part* relevant to the question of whether a cognizer is willing to assert something is--to employ a harmless metaphor--*where* he/she places the relevant claim in his/her cognitive economy. At the very least, we can make the modest point clear that when a claim is in fact reflectively placed so as to be treated by the cognizer as comparable to matters that he/she finds wildly implausible or outright false, then the fact (if it is a fact) that claim is otherwise reliably backed becomes entirely irrelevant.

For my purposes, this means that the fact that a claim was/was not reliably-based, in and of itself, plays no role whatsoever in determining whether or not the problem of philosophical skepticism as I have framed it exists. What we are willing to assert is a function of what we have in fact situated (or would, upon reflection, in fact situate) in the appropriate manner within our cognitive economy. Therefore, what generates the problem of philosophical skepticism is a function of the very same thing. Though in the context of the following quote Bonjour is aiming at a point about the proper analysis of justification, specifically with respect to the question of what we are willing to assert, we can agree with him that "external or objective reliability is not enough to offset subjective irrationality (Bonjour 1980, 20)." So, assuming that the reliabilist is willing to make assertions about the external world, it follows as predicted that they have to recognize their own commitment to *more* than just the claim that such beliefs are *hypothetically* reliably-based and hence hypothetically "really" justified in their originally-preferred

sense. They have to admit that they also *would*, upon reflection, place the claim within their own cognitive economy in such a way that they take it to be properly evidentially backed, in the specific sense that they take it to be the kind of thing they could show to be (likely to be) true. But, of course, as fallibilists they also must admit the opposite of this claim. Hence the paradox.

In his own way, however, Goldman came to recognize how issues of a believer's perspective can affect what Goldman himself takes to be the actual "real" justificatory status of some belief. The recognition is perhaps made most dramatic by a startling concession on Goldman's behalf. He eventually comes to draw a distinction between what he calls *strong* versus *weak* justification. According to Goldman:

A belief is *strongly justified* if and only if it is well formed, in the sense of being formed by means of a process that is truth-conducive in the possible world in which it is produced, or the like.

A belief is *weakly* justified if and only if it is blameless though ill-formed, in the sense of being produced by an unreliable cognitive process which the believer does not believe to be unreliable, and whose unreliability the believer has no available way of determining (Goldman 1988, 56).

It is important to see the gravity of this admission; at least as far as understanding how a reliabilist can use their alleged insight into the concept of justification to respond to the problem of philosophical skepticism is concerned. In admitting, in effect, that we are justified in some positive (even if "weak") sense that a subject's perspective determines, even when the belief is not reliably-based, Goldman in effect admits that how we would go about answering a higher-order worry in an actual argument is relevant to the question of whether or not we are "really" (weakly) justified. Someone who believes something that is only weakly justified behaves rationally in the sense that the claim in question does play the proper role in their cognitive economy: they do not believe something, in

this case, that a little reflection would reveal to be at the very least quite worthy of doubt. Moreover, at least *in certain cases* someone in a weakly justified predicament with respect to some external world claim would be someone who is willing to assert it. In other words, being willing to assert something is in certain cases at least one way to fail to (be positioned to) believe that a belief is not reliably-based, even if it is. Such a person would be one of the blameless sort that would think that there would be *nothing easier* than showing to some challenger that the belief in question was what anyone who was concerned with getting at the truth should believe. Goldman himself, presumably, would think that he was at the very least weakly justified (if not strongly justified) in maintaining things about the external world. But by acknowledging that there is something called weak justification, Goldman only emphasizes how he takes himself to have some subjective-rationality-based grounds for making external world assertions--namely, the fact that doing so is blameless in the specific sense that is appropriate for one who is willing to make assertions about the external world. Importantly, someone who was blameless *in that sense* would obviously *never* say something like "p and I have no basis for ably countering someone who maintains that not p." Someone who was weakly justified in the sense compatible with being willing to make the relevant assertion would never say this, of course, precisely because they would take themselves to be able to counter someone who maintains that not p. So, Goldman's emendation allows him to easily avoid the absurdity of crude reliabilism: i.e., of embracing the *naturalness* of saying things like that. But it just as well indicates how Goldman, assuming he is someone who is willing to make assertions about the external world, should quite readily recognize the existence of the problem of philosophical skepticism. If Goldman takes

himself to be at the very least weakly justified in the sense that a person willing to make assertions about the external world is weakly justified (and hypothetically strongly justified as well) then he takes himself to be able to effectively counter someone who maintains that not *p*. But since Goldman is also a fallibilist, he will also be the first to recognize that no such countering can take place.

Thus, in sum, the developments that Goldman himself deemed to be required in order to continue to uphold reliabilism only serve to emphasize the fact that the problem of philosophical skepticism exists. Again, perhaps Goldman would be unbothered by this fact, given his expressed indifference at times regarding the skeptical problem. And, to repeat, I do not think that the issue I am raising bears on anything other than the reliabilist's ability to respond to the paradox of philosophical skepticism. The overall plausibility of this view does not, of course, hinge on its failure/success in this regard. Nevertheless, given that a paradigmatic reliabilist, when developing his views, runs smack dab into admitting the existence of the sorts of considerations that compose the skeptical paradox, this should show quite clearly why no reliabilist resolution of that paradox will be forthcoming. Quite the contrary, the paradox is actually brought into *greater relief* if the reliabilist is right about justification and also willing to make the relevant assertions.

The odd position that can result from (unlike Alvin Goldman) *consciously* trying to balance reliabilist-friendly views on justification with failing to admit the paradox of skepticism is perhaps most pronounced in some of the writings of Ernest Sosa (specifically Sosa (1993, 1994)). Sosa (as a good fallibilist) admits that we cannot argue effectively for the claim that we possess external world knowledge (Sosa 1994, 96). He

takes this to mean that there cannot be a *legitimizing* account of our knowledge. What Sosa denies, however, is that a non-discursive optimist position (what he calls "formal externalism") will thereby fail to be able to provide us with a general and satisfactory understanding of our knowledge (Sosa 1994, 95). Sosa thinks that a certain kind of circularity (or regress) is inevitable when attempting to attain a general understanding of our knowledge. But, interestingly, he also thinks that because avoiding circularity (or regress) is impossible for "simple, demonstrable logical reasons" we should not think that our inability to do so presents a paradox (Sosa 1994, 109). Instead we should happily accept that merely believing in the reliability of our accepted epistemological practices, and even using those practices to come to believe in that very same sort of reliability, is the best *possible* position for a cognizer like us to be in (as long as, of course, those practices are in fact reliable).

Sosa is quick to admit that factors pertaining to subjective rationality play a central role in attaining knowledge (see, for example, Sosa 1994, 106-107). But he considers it equally correct to maintain that reliabilist-type conditions are also central.

To sum up: We can legitimately and with rational justification arrive at a belief that a certain set of faculties or doxastic practices are those that we employ and are reliable. That remains so, even though someone mad can weave a system of comparable internal coherence and can thereby attain a comparable degree of internal justification [i.e., be equally subjectively irrational]. But in granting this we must not grant that such coherently rational belief need only be true in order to be knowledge. A coherently rational belief can fail to be [reliably-based], surely, and can even be mad if formed by a mind that is really logical though deranged in its social and physical perception and perhaps also in its memory. (A rationally coherent belief *can* also be [reliably-based], of course, and can thereby amount to knowledge as well.) Anyhow, the point remains: there is no obstacle in principle to our conceivably attaining rationally coherent belief in some general account of our epistemic faculties and their reliability (Sosa 1994, 108).

But how can we reconcile the *reasonability* of our coherent general account of our epistemic faculties and their reliability with what Sosa himself admits will be inevitable issues of circularity (or regress)? He addresses the issue as follows:

Perhaps the dissatisfaction...arises from the following reasoning:

If we justify our belief in the reliability of our [way of forming beliefs] W-B:R(W)--by noting that W itself yields B:R(W), then anyone with a rival but self-supporting method W* would be able to attain an equal measure of justification through parallel reasoning. They would justify their beliefs B:R(W*) by noting that W* itself yields B:R(W*). So are we not forced to conclude that someone clever enough could attain a measure of rational justification equal to ours so long as their way of forming beliefs, W*, turned out to be, to the same extent, coherently and comprehensibly self-supporting?

If *this* is the source of the discomfort, then it is discomfort we must learn to tolerate--though in time reason should be able to dispel it...After all, discursive, inferential reasoning is not our only faculty; and logical brilliance does not even ensure sanity. In light of this, I see no sufficient argument why we must settle, at the end of the day, for any irresolvable theoretical frustration...

The desire for a fully general, legitimating, philosophical understanding of all our knowledge is unfulfillable. It is unfulfillable for simple, demonstrable logical reasons. In this respect it is like the desire to find the saint who blesses all and only the nonselfblessed. A trek through the Himalayas may turn up likely prospects each of whom eventually is seen to fall short, until someone in the exhibition reflects that there could not possibly be such a saint, and this for evident, logical reasons. How should they all respond to this result? They may of course be very unhappy to have been taken in by a project now clearly defective, and this may leave them frustrated and unsatisfied. But is it reasonable for them to insist that somehow the objective is still worthy, even if unfortunately it turns out to be incoherent? Is this a sensible response? How would we respond if we found ourselves in that situation? Would it not be a requirement of good sense or even of sanity to put that obviously incoherent project behind us, to just forget about it and to put to our time to better use? And is this not what we must do with regard to the search for fully general, legitimating, philosophical accounts of our knowledge (Sosa 1994 pp. 107, 109)?

I think that if Sosa was right, that the search for a legitimating account of our knowledge was much like the search for a saint who blesses all and only the nonselfblessed, then we could realize that there was no cause for being frustrated by our

inability to explain why we have good evidence for our external world beliefs after pursuing in earnest a first person, neutral examination of the matter. But the analogy is weak in a crucial sense: presumably most people never thought strongly *to begin with* that there was, or had to be, such a saint. In contrast, most people *are* continuously willing to make assertions about the external world, which means that they think that we continuously can back up the relevant claims. This is *all and only* what makes it so that there is a problem of philosophical skepticism. So, a better analogy to draw here would be if we learned for simple, demonstrable logical reasons that $2 + 2 = 5$. In this case the logical reasons would show to be false what we thought was obviously true (and even *continue* to think is true after hearing those same reasons).

Apparently, then, Sosa does not see us to contingently be people who think that we can back up the claims that we make about the external world in the "legitimizing" sort of way. But, in a manner that I find utterly befuddling, he *does* maintain that we are people who think sense perception is reliable (Sosa 1994, 106). Now, what would it be to be someone who thinks that sense perception *is* reliable but who *doesn't* think that they are able to show someone who doubted the truth of a perceptually-based claim why they are incorrect? Wouldn't one's belief in the reliability of perception cause one to proceed confidently in attempting to show a doubter why they were incorrect by placing that doubter within perceptual contact of the relevant truth conditions? Wouldn't we, believing what we do about perceptual experience, think something like: "Surely he will believe me after he sees it for himself!"? This seems beyond obvious. Therefore, it is extremely difficult for me to understand why Sosa can think that discovering for simple, demonstrable logical reasons that there can't be a legitimating account of our knowledge

is like discovering in the same way that there cannot be a saint who blesses all and only the nonselfblessed. For anyone who believes that sense perception is reliable, it seems clear that the more analogous situation would be discovering simple, demonstrable logical reasons for thinking that $2 + 2 = 5$. As far as I can tell, the only way we could reconcile matters in a way favorable to Sosa's position that there is no paradox here, is by saying *not* that people believe that sense perception is reliable, but rather that they believe that sense perception is either *hypothetically* reliable or else reliable *for-all-practical-purposes*. Some of the other writings of Sosa suggest that the latter is what he *also* believes (see, for example, Sosa (1993, 216-218)). In this case it would be true, I think, that people could learn to rationally live with the fact that they couldn't argue effectively for their external world beliefs because of simple, demonstrable logical reasons. But, as I've already shown, this is not the situation of those who are willing to make assertions about the external world.

Here again, we see that my account depends upon the truth of the contingent fact that human beings are willing to assert things about the external world, in such a way that they contextually imply that they can show someone who doubted the truth of the claim (or, as the case may be, the truth of the claim that the person had strong enough evidence for the claim) why that person is wrong. But notice that Sosa seems forced to agree with me. It is quite obvious that anyone who believes sense perception is reliable is one who thinks that they can sometimes back up their assertions in this sense. Sosa's concession on this point, therefore, makes his overall view *extremely* difficult to understand. To me, it points to the oddity of non-discursive optimist positions in general.

3.5 Epistemological Naturalism

I turn now to the final metaepistemology I will examine in this critically-focused chapter; that of the naturalized epistemologist. In examining how an epistemologist friendly to naturalism would respond to the skeptic, I will rely exclusively on the work of Quine, and what I see to be helpful development of that work vis-à-vis the naturalized epistemologist's response to skepticism.⁴⁴

A naturalized epistemologist shares the view of the reliabilist, that what goes into being a justified belief and/or case of knowledge can be explained in ultimately naturalistic terms. However, an important difference is that only reliabilists (or at least reliabilists like Alvin Goldman) believe that they can employ a naturalistically-based account as a means for answering *a priori*-based meta-epistemological questions about, for example, what it means to know, justifiably believe, etc. The naturalized epistemologist, in contrast, rejects the claim that there exists a substantive distinction between the sorts of questions that philosophers as opposed to scientists either actually ask or should ask. He therefore rejects that there exists a tenable *a priori*/*a posteriori* knowledge distinction. By rejecting the notion of *a priori* knowledge, he thus in turn rejects the idea that we can come to know of anything that holds of *necessity*. Everything we believe, according to the naturalized epistemologist, is subject to revision at some point or other as a result of epistemological considerations. It is in this precise sense that he is a fallibilist.

So far, of course, the naturalized epistemologist has said nothing inconsistent with pessimistic fallibilism. In order to understand where these two positions differ, we must understand what makes a naturalized epistemologist a non-discursive optimist. Before

doing so, however, I must clear up the following terminological issue. To "naturalize" the understanding of something is sometimes taken to mean "state what constitutes the matter solely in terms of what is, at least in principle, *able to be studied by science*." Thus, one who would expect me to explain how a naturalized epistemologist "naturalized" knowledge, according to this understanding of that term, would expect me to explain how knowledge is something that is constituted solely in terms of scientifically respectable phenomena. But, for the reasons we have already discussed that have to do with the problems with performing conceptual analyses, it is beyond reasonable to think that no such understanding of knowledge will be forthcoming. Equally clear, however, is that it is an obvious confusion to take the naturalized epistemologist's inability to do so as a criticism of their position. No naturalized epistemologist would consider it at all worthwhile to even attempt to explain what "constitutes" knowledge (or anything else for that matter); where, specifically, to do so means perform a kind of conceptual analysis that is supposed to be free of conceivable intuitive counter-examples. Importantly, however, we also know at this point that to reject this project is *not* tantamount to simply assuming a naturalized epistemologist's view from the get-go. The pessimistic fallibilist agrees that there is, at least as far as we can tell, little point to attempting such analyses.

Where the naturalized epistemologist will differ from the pessimistic fallibilist is primarily in terms of the choice of *methodology* for arriving at claims to know. Whereas the pessimistic fallibilist wishes to enact a first-person neutral methodology, the naturalized epistemologist will, instead, opt for the following sort of inquiry: whether we can call a claim knowledge will be understood to be a function of the degree to which it is an inextricable part of a larger *theory* that itself conforms well with accepted *observation*

sentences, results in the fruitful use of the *hypothetico-deductive method* for making predictions based upon such theories, and has an architecture that results from being guided by considerations having to do with its overall *simplicity, consistency with other accepted, well-entrenched theories*, as well as (perhaps) conformity with other so-called "theoretical virtues." In order to frontload what I think is the most crucial feature of the naturalized epistemologist's response to the problem of philosophical skepticism, it is not necessary for me to explain precisely what it means to honor such constraints. Instead, I can get to the heart of the matter by now addressing the following *meta-methodological* issue. According to the naturalized epistemologist, the very source of understanding that knowledge is properly construed in the above sense is nothing other than the implications of currently accepted theories which, themselves, exemplify the above characteristics. Let us call any theory which exemplifies those characteristics "good science." (It should be noted, however, that the term is imperfect in that it is supposed to include under its extension ways of forming beliefs employed by even putatively non-scientifically-minded folks.) With this, we can sum up the crucial meta-methodological point as follows: according to the naturalized epistemologist, good science itself tells us (or at least can tell us) what good science is, and hence what thereby constitutes knowledge. So, according to a naturalized epistemologist we can call "knowledge" nothing other than "what good science accepts as true;" and we can also say that included in what good science accepts as true is what features a claim must possess in order to be something that good science accepts as true.

The knee-jerk reaction that philosophers have here is to call this understanding of knowledge problematically circular. What the naturalized epistemologist wants to call

knowledge is best seen, according to this common reaction, as rather knowledge *relative to the views of good science*. Therefore, the argument will run, it only amounts to knowledge *if* good science in fact gives us knowledge. And, importantly, at the very least you cannot *argue effectively* for the claim that what good science *says* is knowledge constitutes knowledge by appealing strictly to premises that state what good science says on the matter. To do so is, as it is sometimes put, *epistemically circular*: the premises are justified only if we accept the conclusion they are themselves meant to support. For this reason, the naturalistic way of dealing with skepticism is often rejected.

But this objection is simply wrong-headed. In fact, realizing why it is points directly to the way that a naturalized epistemologist adopts a form of non-discursive optimism. To dismiss the relevance of naturalized epistemology for addressing skepticism as a result of this kind of circularity is, I think, to completely miss the implications of rejecting the a priori/a posteriori knowledge distinction. The naturalized epistemologist can rightly claim that without that distinction, there is little left to motivate the position that there is *anything but* knowledge according to what I have called "good science." To reject a priori knowledge, after all, is to reject the claim that as philosophers we can achieve a standpoint outside of our already contingently adopted, everyday epistemological practices so as to make a fair "unbiased" judgment regarding whether what we judge antecedently to be reasonably arrived at truths are really reasonably arrived at truths. To reject a priori knowledge, then, is to realize, first of all, the *practical* point that if we don't allow ourselves the epistemological practices contingently provided by good science then we will have no means *whatsoever* for arriving at reasonably held truths. This is not to say that we can't or shouldn't examine whether what we take to be

good science really is good science; of course we should. The practical point is, rather, that in order to do so we *must* employ the practices provided by good science--if only because we have no other choice. But furthermore, it is crucial to see just as clearly the *epistemological* point that since there is no such thing as a priori knowledge (at least according to the naturalized epistemologist) it is wrong to think that our being forced in this way reflects some inescapable shortcoming in this regard--for example, that we are doomed to arbitrariness or relativism.

In this way, the naturalized epistemologist agrees it seems, at least in a certain respect, with Wittgenstein: the only sort of knowledge there is, is knowledge *according to a certain world picture*; and it is a confusion to think that the fact that it is knowledge relative to that world picture somehow prevents it from being "really" knowledge. However, of course the views of Wittgenstein and a naturalized epistemologist diverge in many ways. For one, a naturalized epistemologist has a very specific picture in mind as what contingently dictates matters, namely what good science tells us. Secondly, a naturalized epistemologist has a very specific motivation for saying that the latter constitutes knowledge, namely what follows from the failure for there to be anything called a priori knowledge. Lastly, quite unlike Wittgenstein, a naturalized epistemologist would consider it just another part of ordinary epistemological practice to seek to explain why the scientific world picture is itself justified. To perform the latter task is, after all, just what it means to practice naturalized epistemology.

Thus, in sum, what the objector thinks is knowledge *relative to science* is best seen as knowledge *simpliciter*, where the epistemological practices that contingently determine whether something is knowledge are also able (after, that is, we realize the

implications of the fact that there is no such thing as a priori knowledge) to be employed so as to *understand that* they themselves do, *in fact*, determine whether something is knowledge. The only hope for the above objector, then, is to provide good reasons for why there is an a priori/a posteriori knowledge distinction. But, in any case, for the purposes of the present examination we can grant that no such reasons exist; after all, the pessimistic fallibilist thinks as much. With this, at least for my purposes, the worry over circularity can be put to the side.

So, what makes a naturalized epistemologist a non-discursive optimist? To ask this is to ask, in effect, what allows a naturalized epistemologist to be a fallibilist who both rejects discursive optimism as well as uphold what he/she deems to be the proper understanding of the fact that we know quite a lot about external reality? Well, as far as fallibilism is concerned, as I have already pointed out the naturalized epistemologist maintains that we are not absolutely certain of anything at all, including the truths of good science. This view falls directly out of the rejection of the a priori/a posteriori knowledge distinction. The naturalized epistemologist claims that rejecting that there is a standpoint outside of good science is tantamount to rejecting that there is a *completely secure* standpoint that allows us to attain such certainty. Instead we in fact assume a standpoint that causes us to continue to believe/no longer believe claims based upon the end result of a process of weighing theoretical considerations of the sort mentioned above; all the while (however automatically or unconsciously) with an eye towards holding an overall position that allows us to cope with reality (often by the prediction and control of future events). As a result of so going about matters, *any* claim might be given up no matter how obvious it was thought to be beforehand: this even includes claims like

$2 + 2 = 4$ and $p \rightarrow p$. But the naturalized epistemologist also rejects discursive optimism, not *just* because the project of recovering external world knowledge/justified belief in that way is doomed to failure; but *also* because this position wrongly takes fallibilism to imply an epistemic predicament that warrants jumping through such argumentative hoops in order to secure knowledge/justified belief concerning the external world. Though fallibilism may admittedly be a kind of predicament, it should be no surprise that for the naturalized epistemologist--again, in a manner similar to what Wittgenstein claims--what holds for what claims we take to be reasonable holds equally well for claims we take to be unreasonable. Even negative epistemological assessments, that is, must be understood in a manner relative to good science. In this way more *global*, sweeping negative epistemological assessments will obviously have to fail to gain purchase.

But is the latter claim correct? Can't a global negative assessment about our epistemic predicament be reached via a *reductio ad absurdum* of the naturalized perspective? This is an old maneuver. Assume, that is, that what science tells us is correct. It follows that every bit of information we receive about the world is received through the stimulation of our sensory receptors. However, to be *just stimulations* that are often associated with the appropriate external world cause is compatible with being stimulations caused by something else. And, of course, this means that according to good science itself the stimulations we receive every day are *continuously* compatible with false beliefs about the external world.

From here, it might be thought, it will be relatively easy to prove that to believe in the (likely) truth of anything about the external world is presumptuous at best. But, in fact, I think that matters are not so simple. In order for science to defeat itself in the way

envisioned, it must not be just that according to science being just stimulations is compatible with a case where a certain external world belief is false *because of the obtainment of another every day state of affairs*; instead it must be that, according to good science, being just stimulations is recognized to be compatible with the truth of *skeptical scenarios involving systematic deception*. (Or, what amounts to the same thing, it must be both that according to science being just stimulations is compatible with the related belief being false as a result of the obtainment of certain situations consistent with the received laws of nature, *and* that science itself recognizes skeptical scenarios involving systematic deception to be of that nomologically consistent sort.) The reductio could not work even if according to good science the matter of whether or not those stimulations were compatible with the skeptical scenarios was *left open*. In the end, I think this is because you simply cannot establish *fallibilism* without establishing that the skeptical scenarios in particular are problematic epistemic possibilities (see chapter 1 section 1.3). However, I will not be able to make this claim sufficiently evident until chapter 6. But even if I am wrong on this last point, it is clear enough that without it being established or accepted that the skeptical scenarios are epistemically possible, the usually offered *motivation* for skepticism (namely, the dream argument) loses at least some of its force. But for the same reason establishing or accepting the *success* of the attempted reductio here depends on its being true that science *recognizes* or at least can recognize the epistemic possibility of the skeptical scenarios.

But proper examination shows that science cannot do so. To see this, let us grant what would be the ideal case for the one who wants to offer this reductio. Imagine that some scientists discover that they can construct a world for a human subject, a world

subjectively indistinguishable from their own, while keeping the subject unaware that they are doing so the whole time. They would have to transplant an embryo/infant before it was even aware in the womb, and have it exist somewhere else, fully embodied, though in a coma-like state, until, let us say, well into its teenage years, constructing the subject's "dreamt" reality the whole time. They could then wake him/her up and ask "How'd we do?" Let us say that the subject, understandably a little groggy, says after some prodding "Well...not bad at all actually. In fact, at the moment I can say that I cannot, in a general sense, distinguish the way things are here from the way they were before." Let us also suppose that the subject is right, and that the scientists (somehow) learn of this. Now, the thought, of course, is that this probably should at least cause a collective chill to run up the scientists' spines.

But still, notice that *accepting* that their results applied to their own case, even in some remotely worrisome epistemological and metaphysical sense, could only happen through their accepting the conclusion of an argument from analogy. They would have to conclude: "Since we *can clearly do it to unwitting subjects, it could*, in a similar fashion, be happening to us right now!" What I would now like to argue, however, is that it would be incoherent for the scientists to treat this argument from analogy as anything other than weak. This conclusion, if sound, will show that the *reductio* fails; that science cannot coherently recognize the epistemic possibility of the skeptical scenarios.

The reason my desired conclusion follows is this: if the scientists were to treat the argument from analogy as persuasive, and thus treat their discovery *epistemologically* seriously as showing that there is post-discovery, for them, an even *very slight* possibility that they have been, say, dreaming all along, then this just as much undermine the claim

that they had just *discovered* that an unwitting subject's reality *can* be constructed. But if they have undermined the claim that they have just *discovered* this, they have thereby undermined the very basis for their worry over their own epistemic situation. So, to make the presence of the relevant sort of incoherence evident, we can sum up the point as follows: if they discover a scientifically respectable reason to doubt the veracity of their own senses--and this is the best case scenario, it seems, for anyone who wants to go on to offer the above reductio--then it follows that they have *not* discovered anything of the sort. Thus, they cannot discover anything of the sort. They *should* feel the chill, indeed, but there is no good reason to treat the worry that they might be dreaming as anything more than a worry over what is merely an *intelligible* but quite fanciful scenario; the kind of fictional scenario that people have made movies about. But as soon as they take it to be *more* than that as a result of what they have discovered, they undermine the idea that they have discovered that it can happen, and thus undermine their very reasons for taking it to be more than that. Their argument from analogy must be weak.

The naturalized epistemologist has thus evidently emerged quite strong as a result of his ability to deflect such criticisms. We have seen that both the circularity charge (barring, anyway, a subsequently offered plausible argument for the existence of a priori knowledge) and the reductio discussed above are quite without merit.

But, now that the reader has a feel for the naturalized epistemologist's position, what is important to examine is whether a naturalized epistemologist can effectively deal with the problem of philosophical skepticism as I have framed it. In this regard, we have seen that the way a naturalized epistemologist deals with the pessimistic fallibilist is by attacking the very coherence of the idea that his first-person, neutral methodology can

unearth substantive epistemological claims. The naturalized epistemologist thinks that the pessimistic fallibilist is confused if, in arriving at his epistemological conclusions, he refuses to make use of scientific results. Doubly so, in fact, given that the pessimistic fallibilist confusedly thinks he *has to avoid* the use of scientific results in order to legitimate the presence of external world knowledge (and thus, perhaps to a limited extent, thereby legitimate science itself).

However, although this is indeed what a naturalized epistemologist would try to do to attack the pessimistic fallibilist, it should be clear enough that what he lacks is a sound basis for such an attack. We have seen that the basis *usually* given as the reason why we should switch to the naturalized epistemologist's methodology is the failure for there to be an a priori/a posteriori knowledge distinction. It is the fact that no such distinction exists which, we have seen, is supposed to allow us to see the truth in the naturalized epistemologist's particular brand of fallibilism; the very brand of fallibilism that, it appears, allows him to ably uphold a form of non-discursive optimism. But, of course, the pessimistic fallibilist never believed strongly in the claim that there was anything called a priori knowledge. Now, importantly, this does *not* show that there is an *uncertain connection* between the claim that there is, in fact, no a priori knowledge and the claim that the naturalized epistemologist's methodology steers us aright. We can grant the naturalized epistemologist, at least for the sake of argument, the claim that there is such a strong connection. We can grant him, that is, that if he is right that there is no such thing as a priori knowledge, then his methodology is unequivocally the correct one--and that he is able to deal with the circularity issue as well as the above attempted reductio. We must recall, however, the pessimistic fallibilist's reason for failing to countenance an

a priori/a posteriori knowledge distinction. The reason was out of a desire to remain *neutral* on controversial matters. This means that the pessimistic fallibilist sees the claim that there is no a priori knowledge to be *just as controversial* as the claim that there is. The pessimistic fallibilist thereby assumes a perspective that sees the following to be the inevitable outcome of any argument one way or the other with respect to this issue: he sees it as inevitable, that is, that the proponent of a priori knowledge will, when presenting a valid argument for their position, beg the question against the challenger of a priori knowledge, *and vice versa*. Notice that in this case the naturalized epistemologist's usually unapologetic reference to good science will be of no avail. The latter maneuver only works when it is taken as granted *as part of good science* that there is no such thing as a priori knowledge; that the human organism gathers well-founded beliefs about the mind-independent world solely through his perceptual faculties. To state the point differently, the latter maneuver of the naturalized epistemologist only works when we understand what falls quite directly out of the sort of world picture which treats the methodology prescribed by good science as contingently how we go about our epistemological business (including when we do naturalized epistemology). However, if, in contrast, it is *not* taken as obvious that there is no a priori knowledge, then we have no basis whatsoever for unapologetically employing the naturalized epistemologist's methodology. The pessimistic fallibilist adds to this last conditional claim, the equally plausible conditional claim that when we start from *neutral ground* that begins by treating it as not immediately obvious one way or the other whether there is a priori knowledge, that there will be no way to resolve the relevant dispute in a way that shows either of the opposing views to be the clear winner.

It is important to see that the naturalized epistemologist's previously-discussed maneuver for explaining why any apparent circularity in any argument they offer is unproblematic, will not work to explain why they can again unapologetically beg the above question. They were quite allowed to so proceed before because of what seem to be clear enough implications of the failure for there to be a priori knowledge. When it is no longer granted as obvious that there is no a priori knowledge, the naturalized epistemologist is no longer allowed to get away with doing so. Everything for the naturalized epistemologist depends upon the obviousness of the claim that there is no standpoint outside of science from which to judge whether science is correct. Given his desire for neutrality regarding controversial theoretical matters, the pessimistic fallibilist would avoid accepting the obviousness of this claim.

I will not review yet again why the pessimistic fallibilist is still, even given his extremely neutral standpoint, able to indicate that a quite compelling problem of philosophical skepticism exists. But, it is for this simple reason that the naturalized epistemologist fails to deal with that issue. In sum, the only way the naturalized epistemologist evades the paradox is by being dogmatic in the eyes of the pessimistic fallibilist, and his usually unapologetic deference to good science as an explanation for why such circularity is okay will, for the already discussed reasons, be ineffective.

3.6 Conclusion

This concludes the critical discussion of non-discursive optimism. We have seen that each of the above positions discussed fail to evade the problem of philosophical skepticism. That problem arises as a result of a mismatch between what we contingently take our epistemic position to be with respect to external world claims, and what

fallibilism entails regarding that same epistemic position. What we have seen is that these two conflicting factors are quite fundamental to a majority of the non-discursive optimist positions discussed above. The contextualist and process-reliabilist must acknowledge the presence of the two conflicting factors as a result of what seem like necessary developments of their own positions; while Wittgenstein must do so after we merely adjust the terminology appropriately so as to be able state the essence of the issue in language he would accept. Only the naturalized epistemologist can, it seems, maintain a position that more or less easily fails to acknowledge the presence of the above factors, but she does so at the price of requiring dogmatism in her belief that there is in fact no a priori knowledge--in other words, that there is no standpoint outside of good science.

With this we also conclude the more expository part of the work. The task now will be to take what I think are the necessary steps for the proper resolution of the problem of philosophical skepticism. What the arguments so far have hopefully established is that no such resolution can take place if it is granted that fallibilism is correct. And, to be sure, in the next chapter I will attempt to cast doubt on the truth of fallibilism. But this is only the first step in the resolution. As we shall see, an even stronger form of the same kind of "mismatch" problem emerges if all that is established is that from a first-person, neutral standpoint there is no convincing case for fallibilism. What is needed in order to resolve the problem once and for all is an explanation of how from that same standpoint a convincing argument *for infallibilism* can be mounted. I will reserve that positive project for chapter 5. Chapter 6 will then contain helpful developments of the positive meta-epistemology/theory of perception that I present.

CHAPTER IV DOUBTS ABOUT FALLIBILISM

4.1 Introduction

Recall the Dream Argument. In that argument a crucial inference was made from the intelligibility of the skeptical scenarios to the claim that it is consistent for us to receive all of our neurological deliverances while nonetheless having beliefs about the external world that fail to be non-accidentally true. In this chapter the soundness of that inference will be questioned. In order to do so I will critically examine a rather persuasive argument for fallibilism that has been proposed by Barry Stroud.

Stroud continues to convincingly defend what he calls the *conditional correctness* of fallibilism.⁴⁵ In his view, once we grant that we can legitimately undergo an "external" philosophical examination of our knowledge, the fallibilist conclusion becomes inevitable. An "external" examination is here understood to be a *detached*⁴⁶ *assessment of the totality of our knowledge of mind-independent reality.*⁴⁷ This can be seen as, in effect, Stroud's preferred way of understanding what I've been calling a first-person, neutral methodology. He sees activity like this to be paradigmatic, Cartesian-style philosophy; a kind of broadly focused "anthropology."⁴⁸ Doing so, he argues, makes clear that:

- (1) It is a necessary condition for my knowing anything about the external world that I know I am not dreaming.

The heart of Stroud's position seems to be that (1)'s utter obviousness⁴⁹ makes defending the position that we know anything about external reality *impossible* in the following sense: we cannot consistently both grant (1) while engaging in an "external" examination as well as adequately explain why we have any such knowledge. In order to explain why we have it, we must explain why we know, or at least are able to come to know, that we

are not dreaming. However, our only tool available for establishing that we are not dreaming is perceptual experience, and in Stroud's considered view--one that takes off from consideration raised by DA--perceptual experience is simply not up to the task.⁵⁰

Part of appreciating the power of the skeptic's position involves clearly indicating how and why particular responses to it seem to fail. Stroud argues that, in certain cases, it is the way that an otherwise plausible-sounding response violates the requirements⁵¹ for being an "external" examination, which explains why it misses the mark. Fittingly, at times he calls this sort of unsatisfying answer an "internal" response to skepticism. The purest example is, without question, the Moorean response (Moore 1939). I think Stroud is quite right in maintaining both that Moore's response fails, and that we can learn something important about the problem of skepticism by identifying why it fails. However, in this chapter I will argue that he cannot consistently find Moore's answer inadequate in roughly the way he does, while also persuasively maintaining the conditional correctness of fallibilism. I agree with Stroud that (1) holds generally; and henceforth I will assume that the conception of knowledge that goes along with maintaining this is correct (see n. 49). Nevertheless, from a properly understood "external" standpoint it is inconsistent to disallow Moore's affirmation of the relevant antecedent of (1) while wholeheartedly endorsing, as Stroud does, the denial of the consequent. Since, as we shall see, Moore's unabashed confidence in the reasonability of his perceptually-based claims is indeed misplaced given the kind of "external" question that the philosopher asks, I argue that *so is* Stroud's lack of confidence. A proper understanding of what it means to assess the *totality* of our knowledge is what reveals this parity between Moore and Stroud. As a result it is, at the very least, no longer

compelling to maintain that fallibilism is conditionally correct. No compelling reason exists for thinking that it is the *inevitable* result of a properly understood philosophical examination of our knowledge. Another upshot is that the matter of whether or not fallibilism is warranted must be looked at in a novel way. It requires probing an issue that is more nuanced than a methodologically constrained survey of one's stock intuitions about the powers of perceptual experience.

4.2 Cartesian Questions

As I've already stated, Stroud maintains that a properly philosophical "external" examination will involve a detached assessment of the totality of our knowledge of mind-independent reality. It is not necessary to here critically discuss each of the important characteristics of this kind of assessment.⁵² Given my aim, my critical attention in this section will be centered only on what it means to honor the *totality* condition. This condition is clearly the most problematic for anyone wishing to defend epistemological realism about the external world in a way that is supposed to address the concerns of a Cartesian-style skeptic. This is because the totality condition has us assess *all* of what we take to be our knowledge of mind-independent reality, *all at once*. The reason this otherwise rather draconian-sounding condition is in place is simple. If in a first-person, neutral examination we are after an explanation of why *anything* we take to be external world knowledge is (in certain cases at least) really knowledge, then it does no good to rely on merely presumed instances in order to make our case. Interestingly, however, Stroud develops this clear enough meta-methodological consequence in the following way:

If we start by considering a certain domain of facts or truths and ask how anyone could come to know anything at all in that domain, it will seem that any other knowledge that might be relevant could not be allowed to amount to already knowing something in the domain in question. Knowledge of anything at all in that domain is what we want to explain, and if we simply assume from the outset that the person has already got some of that knowledge we will not be explaining all of it (2000, 103).

Here we can see that, in Stroud's view, it is not just that we cannot rely on any presumed piece of external world knowledge in order to make our case; it is also true that we cannot rely on *anything* that amounted "to already knowing something in the domain in question." However, to say that a first-person, neutral examination of our knowledge of external reality requires this is, as I will now begin to illustrate, simply false.

The wording of the above passage strongly suggests that Stroud incorrectly sees violating the just-stated requirement as tantamount to *simply assuming* "from the outset that the person has already got some of that knowledge." Of course, it is true enough that one hasn't rendered legitimate the *totality* of whatever it is proper to call one's "knowledge of the external world," if one first merely grants oneself some piece of external world knowledge and then, for example, validly infers more external world knowledge from it. Thus, the only allowable way to proceed is to rely exclusively on knowledge that is not in the domain in question. Certain rationalist maneuvers notwithstanding, a familiar choice here is to rely on knowledge of our experiential states. But, of course, we know from the results of chapter 1 that the project of recovering external world knowledge from this type of knowledge is a hopeless enterprise.

But Stroud evidently interprets the force of the totality condition in a manner stricter than just what can lead to this familiar problematic. In addition to the perfectly reasonable requirement that we not rely on a presumed piece of external world

knowledge, Stroud interprets the totality condition as barring the reliance on *any* sort of knowledge that entails that we know something about the external world. Here is how the above excerpt continues:

Any knowledge we do grant to the person will be of use to him only if he can somehow get from that knowledge to some knowledge in the domain in question. Some inference or transition would therefore appear to be needed -- for example, some way of going from what he is aware of in perception to knowledge of the facts he claims to know. But any such inference will be a good one, and will lead the person to knowledge, only if it is based on something the person also knows or has some reason to believe...That "something" that he needs to know cannot simply be part of his evidential base, since it has to get him beyond that base. But it cannot go so far beyond that base as to imply something already in the domain in question either, since the knowledge of anything at all in that domain is just what we are trying to explain. So it would seem that on either possibility we cannot explain with the proper generality how the kind of knowledge we want to understand is possible (Stroud *ibid.*).

According to this view, the totality condition leads to the following dilemma. One cannot, it seems, both *legitimately* cite a piece of possessed knowledge outside the target domain and also effectively *recover* knowledge within the target domain. If one cites knowledge which is both i) not external world knowledge; as well as ii) *known to entail* that one has some external world knowledge--henceforth, let us call knowledge of this last dual sort *basic entailing knowledge*--then this will amount to violating the totality condition. One would be explaining why one has external world knowledge on the basis of knowledge--namely, basic entailing knowledge--that guaranteed one *already* had it. Basic entailing knowledge recovers external world knowledge, if you will, but (according to this view) only illegitimately. If, however, one tries to keep things legit by citing a piece of knowledge that is known to fail, by itself, to guarantee one already has external world knowledge--let us call this *basic non-entailing knowledge*--then one won't be able to *recover* external world knowledge via that sort of explanation.⁵³ One would not, in this

case, get what one was after: an explanation of why one knows certain things about the external world. At best (and even this result seems quite dubious) one would be explaining why *it is likely* that one knows certain things about the external world. In short, according to Stroud, as a result of the totality condition the boldest of anti-skeptics who accepts the legitimacy of the Cartesian enterprise will, it appears, have to rely on claims that are either question-begging (in two different possible ways)⁵⁴ or else too weak to give the desired result. I will henceforth call this the *realist's dilemma*.

But notice the important gloss that is involved in the above formulation of the realist's dilemma. In what circumstance would it be disallowed to rely on basic entailing knowledge? The answer can only be, it seems, when either the claim that one possessed the more basic knowledge was itself unsupported, or else was supported inadequately (presumably because the argument offered *for it* begged the question). More accurately, then, contra-Stroud one violates the totality condition in the relevant sense when one relies *without adequate argument* on basic entailing knowledge. In the above excerpt, Stroud seems to tacitly conflate relying on basic entailing knowledge with relying on external world knowledge *simpliciter*. Despite what Stroud claims, and in stark contrast with relying on external world knowledge, we *can* explain, with the proper amount of generality, why we know what we do about the external world by relying on basic entailing knowledge. We can do so as long as we are also able to argue in a non-question-begging manner for the claim that we possess the more basic knowledge (it is assumed here that the fact that the latter entailed that we had some external world knowledge would be at least less controversial, if not obvious). The analogous situation does not hold when one relies on a piece of external world knowledge. It is, of course, correct that

relying on *adequately supported* external world knowledge would explain why we know things about the external world just as effectively as would relying on adequately supported basic entailing knowledge. But the brute methodological force of the totality condition lies in how it places constraints on what can constitute adequately supported external world knowledge. We must *start* from knowledge that is *not* external world knowledge in order to adequately explain why we have the latter. But it is equally true that *if* we can adequately argue for the claim that we have the right sort of basic knowledge, then we can recover external world knowledge in the proper way.

We can probe the same issue a bit more deeply by now asking: Is it *possible* to adequately argue for basic entailing knowledge? Stroud apparently thinks not, though in an importantly qualified way. He seems to construe the heart of the problem for the epistemological realist who accepts the legitimacy of a first person, neutral examination but nonetheless tries to offer a constructive account to be this: only arguments that employ sound reasoning can serve to rebut a skeptic, however, any argument offered will inevitably beg the question in some way or other. Therefore, in order to address our question, we should now have a closer look at this all-important worry over question-begging. Doing so will shed light on the precise extent of the trouble that the totality condition causes for one wishing to philosophically establish and defend epistemological realism about the external world. As it turns out, matters are much less dire for the realist than popular philosophical opinion suggests.

The fundamental realization is perhaps this. Generally speaking, a claim regarding what does versus does not beg the question in a certain dispute only makes sense when set against a background that involves, in part, knowledge of the *shared* assumptions

between the disputants. For example, it is common to defend epistemological direct realism by pointing out that it is the only view, among its competitors, which *avoids* skepticism.⁵⁵ In the context of a debate with a representational realist, then, the direct realist is perfectly allowed to use the latter as a virtue of his/her position. Thus, we are right to ask, what presuppositions, if any, might the boldest of realists and the fallibilist share, and how does this clarify what does versus doesn't consequently count as question-begging?

Obviously, as already discussed, the realist and the fallibilist should agree as a result of the imposition of the totality condition that no constructive account can *merely assume*, at any point in the explanation they offer, that they already have some piece of external world knowledge. This means that no constructive account can rely without adequate argument on either a claim to know something about the external world, or else a claim to know some piece of basic entailing knowledge. Note, however, what is versus is not implied by being forced to honor this perfectly reasonable constraint. Note, more particularly, the scope of the "not" in the stricture "one *cannot* assume that one already has external world knowledge." Clearly enough, honoring this is far different from honoring a requirement that gives the term "not" much narrower scope. According to this much stronger stricture, one must assume for the examination in question that one *does not* know anything about the external world--which amounts, in effect, to the requirement that one assume that one *only* possesses basic *non*-entailing knowledge.

It is clear enough that according to Stroud the latter requirement is the one that the skeptic imposes. This would explain why he understands relying on *any sort* of basic

entailing knowledge as equivalent to merely assuming that one already knows something about the external world. As Stroud writes elsewhere:

[(Narrow)] The demand for completely general understanding of knowledge in a certain domain requires that we *see ourselves at the outset as not knowing* anything in that domain and then coming to have such knowledge on the basis of some independent and in that sense prior knowledge or experience. (Stroud 1989, 144, my emphasis).

But, of course, there is all the difference in the world between honoring (Narrow) and honoring the much weaker requirement that one *not merely assume that one already has external world knowledge*. Let us call the latter requirement (*Wide*); so named for the wider scope of the term "not" found within it. Notice that honoring it is compatible with the *additional* requirement that at no point in one's assessment can one rely *without adequate argument* on the claim that one *only* possesses basic non-entailing knowledge--which would amount to the mere assumption that one *already fails* to know anything about the external world. Put in a clearer way, honoring (*Wide*) is compatible with also requiring that one *not merely assume that one does not know anything about the external world*. Let us call the conjunction of (*Wide*) and this last-mentioned compatible negative requirement "(*Wide+*).". In contrast with (*Wide+*), when "see ourselves" is taken to mean "assume," honoring (Narrow) requires that one do just what I have most recently italicized.

This difference is important to recognize if what one is after is an understanding of the extent of the trouble that the totality condition causes for the boldest of realists. Having to honor (Narrow), after all, does make the task of legitimately recovering external world knowledge look to be an impossible enterprise, quite obviously subject to the realist's dilemma. If at the outset of my philosophical assessment I am forced to

assume that I *do not* in fact know anything about external reality, then it is easy to see that the only way I can pull the realist rabbit out of the hat is by surreptitiously placing it there at some later point; i.e., by eventually saying something that is dogmatic or question-begging. But if, in contrast, all we are required to assume at the outset is what is dictated by (Wide+), then we are *forced* to remain open to the possibility of legitimately recovering external world knowledge *via adequately supported basic entailing knowledge*. In a sense, whether or not the constructive project at hand is legitimate would depend on whether the realist rabbit is hiding somewhere in the hat to begin with. The difference between (Narrow) and (Wide+), to speak less playfully, is that only (Narrow) requires assuming a perspective that already makes explaining the presence of common sense knowledge *methodologically impossible*. In contrast, the weaker constraint (Wide+) allows one to maintain, as a matter of methodology, a perspective defined by being neutral with respect to the question of whether or not common sense knowledge might be recovered via adequately supported basic entailing knowledge. It thus goes without saying, contra-Stroud, that (Wide+) is the more appropriate constraint for the "external" philosophical examination at hand. Honoring (Wide+), unlike (Narrow), allows for fallibilism to be an *interesting* result. To make proper sense of how the problem of philosophical skepticism is supposed to fall out of a first-person, neutral methodology, one must understand the totality condition accordingly.

Prominent critical discussions of Stroud's views seem to miss this point. For example, in his rightfully popular book on skepticism Michael Williams concedes, incorrectly I think, that fallibilism will be *inevitable* when we undergo a first-person, neutral examination of the extent of our external world knowledge.

...[T]he source of the doctrine of the priority of experiential knowledge is...the distinctively philosophical project of trying to understand how it is possible for us to know anything whatsoever about the external world. The totality condition that the sceptic (or the traditional philosopher) imposes on a philosophical understanding of our knowledge of the world is what forces us to see the knowledge as somehow derivative from experience...

....[T]his argument...shows...that the doctrine of the priority of experiential knowledge over knowledge of the world is a *methodological necessity of the traditional epistemological project* (1996, 126-127, his emphasis).

To say that my views on skepticism have been influenced by Williams would be a severe understatement. Nevertheless, here I want to note his concession that *given* the skeptic's "distinctively philosophical" methodology, the conclusion that we ought to treat our sensory evidence in a manner friendly to fallibilism follows *as a matter of methodological necessity*. Williams has thus apparently failed to notice that, for the reasons discussed above, (Wide+) rather than (Narrow) is far and away the more appropriate interpretation of the totality condition. He is not alone. The mistake is implicit, for example, in any view which sees classical foundationalism to be the *inevitable* outcome of accepting the legitimacy of a Cartesian-style examination of our knowledge.⁵⁶

We can now sum up what position honoring (Wide+) leaves us in with respect to the prospects of defending realism in a non-question-begging manner. If we are to defend epistemological realism about the external world, we must do so in a way that initially treats realism and the position that we do not know anything about the external world as equally viable. This amounts to *beginning* one's philosophical examination by assuming a perspective that simultaneously fails to merely grant the truth of either position. (Wide+) requires us, however, to also be quite open to the possibility of going on to adequately support realism via *additional* well-chosen premises that both parties would have to also

accept in order to avoid inconsistency. *The mistaken pessimism so often voiced among epistemologists is the result of wrongly maintaining that once we grant that we have to start from neutral ground, there is consequently no possible way that we can then come to validly infer that the skeptic is wrong that isn't question-begging.* A closer look at the phenomenon of question-begging shows that anyone who sees things in this last pessimistic way is someone who, in fact, has sealed the realist's position off initially and thus made skepticism an uninteresting philosophical result. Although the totality condition rightly disallows the *assumption* that we already have external world knowledge, we must still be open to the possibility that we do *in fact* already have it, *and* that we can argue effectively for this claim. There is simply no immediately compelling reason to think that a non-question-begging argument for realism can't be mounted that involves starting from the neutral ground characterized by interpreting the totality condition in a manner that honors (Wide+). What such an argument might look like (or, contrarily, whether the general problem of question-begging does still indeed stand) will depend, in part, on how to properly characterize the positive philosophical project of assessing the totality of our external world knowledge. If the realist can both obey the rules in place for such an assessment as well as cite adequately supported basic entailing knowledge, then he/she will be successful. An attempt to perform this feat will be made in the next chapter.

4.3 Moore's Internal Answer to Skepticism

Having made an attempt to partially explain what must go into being an "external" examination of our knowledge--which is Stroud's term for a first-person, neutral methodology--I will now address how it helps to reveal what seems wrong with a

Moorean response to the skeptic. This will set things up nicely for the next section, wherein I argue that Stroud's view that fallibilism is conditionally correct looks similarly problematic.

In Moore's famous paper "Proof of an External World" the following infamous demonstration is described. Moore first raises a hand and then makes a certain gesture with it while declaring "Here is a hand." He then does the same exact thing with his other hand. After doing so he declares that he has rigorously proven that there are external objects (Moore 1939, 146).

Most philosophers rightly find something wrong with Moore's argument. At the very least, by itself it seems to be a grossly inadequate way to address the *philosophical* question of whether we know that an external world exists.⁵⁷ The above discussion of what goes into asking that kind of question has the virtue of giving a clear explanation for why this is so. In short, Moore violates the totality condition.⁵⁸ It is indeed a perfectly valid inference to move from knowledge that one has hands to knowledge that external objects exist. However, if what we are after when asking an "external" question about our external world knowledge is whether we have any of it at all, we cannot *fundamentally* rely on an external world knowledge claim in order to make our case. Doing so violates (Wide+).

A helpful way to preliminarily state the problem with Moore's argument is as follows. Granting that the following conditional holds:

(1*) I know I have hands only if I know I'm not dreaming,

in the context of an "external" examination, fundamental reliance on the knowledge that one has hands as a way of discharging the consequent of (1*) via modus ponens is illicit.

The requirement that one honor (Wide+) disallows proceeding as such in order to explain why we, in fact, know that we are not dreaming and hence that the mundane external world exists.

4.4 Against the Conditional Correctness of Skepticism

I think the general lesson we can learn from the apparent failure of Moore's "proof" is this. While involved in an "external" examination we have to probe deeper than our stock intuitions about the powers of perceptual experience, in order to properly engage the question of the extent of our external world knowledge. The task of this section is to motivate this claim in a way that makes clear what is problematic with both Moore's response as well as maintaining the conditional correctness of skepticism.

I take it to be clear enough that a view about the powers of perceptual experience can play a central role in what assessment one arrives at regarding the extent of one's external world knowledge. If nothing else, it is evident that one's metaphysics of perception can crucially inform one's epistemology. To wit: these days it is widely agreed that veridical experience is, in a specific sense, logically complex. To veridically perceive, according to this view, it must be that one is having certain perceptual experiences *and* (more or less independently, depending on the view) that those perceptual experiences also "match up with," or "are satisfied by" what is really the case (see Ch. 1, n. 7). In the first chapter, I called this view *the conjunctive theory of veridical experience*. Obviously the naturalistically-based versions of this sort of view would take radical forms of external world skepticism to be an in principle intolerable result. But in the context of trying to rebut the skeptic "on his own terms," the way that this view would hamper one's constructive project is straightforward. We have seen, that is, that *if this*

view is correct one simply cannot *establish* that one's own perceptions "match up" with what is (likely to be) the case, and thus be positioned to refute the claim that matters are otherwise. Therefore, that it makes rebutting someone who denied that we have any positive basis whatsoever for our external world claims impossible is a result that falls directly out of accepting, for example, the obviousness of some version of a conjunctive theory of veridical experience.

This last realization provides the key to understanding what the inadequacy of Moore's response to skepticism reveals about the fundamental difficulty involved in refuting the fallibilist. Moore's failed attempt is best seen as bringing properly to light the dynamical relationship that exists between what sort of general take on perceptual experience might be common coin for the realist and the fallibilist, and how the prospects for constructively defending realism consequently look. The relationship between granting, for example, *the truth* of a conjunctive theory of veridical experience and the consequent *difficulty* a realist will encounter when nevertheless heroically attempting to rebut the fallibilist, is just one facet of one kind of relationship that might exist between the shared background assumptions of the realist and the fallibilist, and what consequently counts as question-begging. Taking a more bird's eye view of matters, then, we can sum up the relevant predicament for the realist in the following way: *If* a conjunctive theory of veridical experience is granted initially (as it indeed should be) as *not obviously false*, then reliance on one's sense experiences alone is demonstrably not adequate for philosophically explaining why one knows that the antecedent of (1*) holds. According to this conception of matters, the only way Moore can answer the fallibilist involved in a first-person, neutral examination in the way he does is by begging the question *against a*

conjunctive theory of veridical experience. In other words, Moore seems to think that relying on his perceptual experiences alone *is* adequate for knowing that the antecedent of (1*) holds. In the context of a properly framed philosophical examination of the extent of our external world knowledge, this can be described as ignoring the fact that a certain theory of veridical experience might be true which would make it so that this sort of reliance was insufficient. With certain additional provisos,⁵⁹ this fleshes out the precise reason why Moore can be convicted of *merely assuming* that he knows that he has hands, thusly violating the totality condition. That it requires begging the question against a conjunctive theory of veridical experience in order to achieve its sought after result is--again, with certain provisos--the more fleshed out explanation of why Moore's type of demonstration is not an adequate way to address the philosophical question at hand.

But, we should now ask, can a similar issue be raised for someone who instead is *partial* to a conjunctive theory of veridical experience? I think it is quite obvious that it can. In short, anyone who accepts a conjunctive theory of veridical experience as primitively true thereby rules out by fiat a certain sort of contrary view; a view which *rejects*, in particular, the claim that our veridical experience of mind-independent reality is constituted in any way by more or less determinate contents that may or may not, as the case may be, "match up with" or "be satisfied by" what is really the case.

It will take some time to make this fully clear, but the point is that if we are involved in a properly framed philosophical examination, it follows that we cannot simply rule out a theory which *rejects* a conjunctive view from the get-go. Instead, for reasons I will soon discuss, we have to at least be open to the possibility that the latter sort of theory is true. I have chosen to call a position of the sort in question a *non-*

cognitive theory of perception. Understood in the above negative form, it is clearly quite determinable. That said, I will now submit the following sketch of a prototype; further developments of the position will be given later on, in the final two chapters. For the moment, however, all I need is an evident enough foil for a conjunctive theory.

Suppose, then, that perceptual experience was not constituted in any way by any sort of "contents," but instead was most accurately understood as constituted by the exercising of one's sense organs in a manner that allowed one to comport one's body appropriately in one's immediately surrounding environment (see Ch. 1, n. 8). On this view, for example, one of the primary functions of vision is to *measure relative distance*, in a manner that allows one to comport one's body in the consequently appropriate manner (I am not suggesting, of course, that in the end all proponents of a conjunctive theory of veridical experience need deny this last claim). Among other things, vision allows one to avoid the murderer who is coming at you with a knife; to keep your distance (if you can). The important feature of this sort of theory of perception is how it consistently rejects the claim that perceptual experience is constituted, in whole or in part, by the kind of thing that may or may not "match up with" or "be satisfied by" what is really the case. As such it rejects, for example, the claim that perceptual experiences *by themselves* serve as an epistemologically efficacious intermediary between oneself and certain specific claims about mind-independent reality.⁶⁰ The non-cognitive theorist would, of course, be wise to maintain that perceptual experiences play *some* crucial role in one's day-to-day epistemological practices. What the non-cognitive theorist will deny, however, is that perceptual experience plays this role as a result of its "contents" (be they representational contents or something else).⁶¹

For the purposes of the present discussion, the important point is this. If after undergoing the skeptic's favored sort of examination it was discovered that this last sort of theory was far and away the most plausible one to accept, then this would forcefully block one well-trodden road to skepticism. If, for example, in the context of that sort of inquiry I come to know that *just by undergoing my every day or "plain"*⁶² *waking experience* I am undergoing a process that allows me to comport my body helpfully in an independently existing world (including allowing me to find a place to sleep and therefore, sometimes, dream), then I will thereby know that that experience is not logically complex in the specific sense that a conjunctive theorist maintains. I will thereby know to be false any skeptical scenario involving *systematic* deception, whose consistency with my "plain" waking experience depends on construing it as if it were logically complex in the relevant sense. The discovered reasonableness of a non-cognitive theory of perception within the context of the skeptic's favored sort of examination of our knowledge is what could bear this heavy burden. Of course, by getting the skeptical scenarios out of the way in this particular fashion, one will not have yet explained how it is that we come to know particular things about the external world. A lot of philosophical work would remain, specifically regarding uncovering the role that we consequently ought to think perceptual experience *does* play in our epistemological practices. But one *will* have forcefully countered the sorts of worries which give rise to the problem of philosophical skepticism, and hence addressed what seem to be the only sorts of concerns that could undermine one's ability to philosophically legitimate "common sense." The road to a constructive account of our knowledge of external reality,

even when it is paved in the way a skeptic would otherwise prefer, *must* pass through the philosophical theory of perception it is most reasonable to accept.

With this in mind, to help achieve my present aim of showing why laboring under (Wide+) casts serious doubt on the conditional correctness of skepticism, what I would now like to do now is finally point my crosshairs directly at DA in order to cast doubt on the truth of fallibilism. That is, I will now show that the usual arguments for i) the claim that we do not know we are not dreaming; depend on simply granting ii) the philosophical defensibility of a conjunctive theory of veridical experience, *and vice versa*. If this is correct, it follows that anyone who actually tries to defend the claim that the skeptical scenarios are epistemically possible by appealing to a conjunctive theory, or vice versa, will be caught in a very tight circle.⁶³ This means, in turn, that coming to maintain that we do not know we are not dreaming is simply ill-motivated, at least when we are constrained by having to honor (Wide+). As will become clearer below, the typical arguments offered are either demonstrably invalid or else beg the question against a non-cognitive theory of perception; and, to repeat, to rule out the latter theory *by fiat* violates the totality condition in the relevant sense. It amounts to merely assuming at the outset that we do not in fact know anything about the external world--that we *only* possess basic non-entailing knowledge.

Given that a proponent of a non-cognitive theory has the above-stated plausible response to the manner in which skepticism is usually motivated, it is pretty obvious that the defensibility of a conjunctive theory depends on the defensibility of the claim that at any given time I do not, in fact, know I am not dreaming, a brain in a vat, etc. *If* I had good reason to think that I do not know I am not dreaming, then this would be a rather

forceful reason for rejecting a non-cognitive view, and hence for accepting either a conjunctive theory or else something close enough to it in the important respects. Granting that, illustrating the presence of the relevant problematic circle only requires showing the converse: how defending the claim that the skeptical scenarios are epistemically possible depends on first having adequately defended a conjunctive theory of veridical experience. I will now attempt to show this by explaining how *accepting* the view that the skeptical scenarios are epistemically possible requires accepting, apparently without additional supporting argument, some form of a conjunctive theory of veridical experience.

So, what *are* the reasons one might offer for the claim that neither you nor I know we are not dreaming, a brain in a vat, etc.? As far as I can tell, one would only have the following plausible-sounding claims to present. As I have already pointed out, it is true enough that at any given time we are all able to find intelligible certain skeptical scenarios (I am dreaming, a brain in a vat, etc.) that have the following properties. First, these stories are consistent with one relatively neutral way to describe our sensory evidence--something like "I'm having sense-experiences as if p." (Where by this description being "relatively neutral" I simply mean that it is compatible with the mutually incompatible claims "p" and "I am merely dreaming that p.") Next, it is also true that *were* any of these skeptical scenarios to obtain, then any claim we offered about the external world would turn out either false, or else true merely by accident.

But such grounds are clearly not sufficient for adequately defending the position that I do not know that I am not dreaming, a brain in a vat, etc.⁶⁴ From the fact that the relevant carefully crafted scenarios are both intelligible to state as well as consistent with

one (relatively neutral) way to describe my sensory experience, it simply doesn't follow that I do not know that they fail to obtain. Obviously, the mere fact that the skeptical scenarios are intelligible to state does not imply that they are epistemically possible.⁶⁵ At this point, though, I imagine the following point being pressed: "It is not just *any* intelligible story that is told by the skeptical scenarios, but rather one of a sort that is consistent with a very *neutral* description of my sensory evidence. This alone is epistemologically pertinent because, generally speaking, the most neutral among a set of viable candidate descriptions of our sensory evidence is the one that it is most reasonable for us to accept. It is simply a fact that a description of this sort will always be 'the safest bet.' It will be the description most resilient in the face of skeptical attack. And, whatever its specific character, such a description will clearly be one which is also consistent with the truth of the skeptical scenarios."⁶⁶

This is an admittedly persuasive line of thinking. But it rests on the false principle that, in general, being a more neutral claim implies being a more reasonable claim in the relevant sense. If I'm staring at a dog, e.g., and I make the more neutral claim (perhaps with an accompanying nod) "There's a mammal", it just doesn't follow that I am not, in that case, just as warranted to assert the less neutral claim "There's a dog." Suppose the epistemic situation was such that I initially came to *know* that there was a dog before me, and then inferred from that knowledge *alone* that there was a mammal before me. In that case the two beliefs would obviously have an identical degree of justification. In fact, as matters played out, the sound epistemic status of my more neutral belief would even depend on the sound epistemic status of my *less* neutral belief. It is true that there might be a similar case where I, after being challenged, retract

my claim to know that there was a dog before me but continue to reasonably maintain that there was a mammal before me (Perhaps I was staring at an animal that looked an awful lot like a wolf). But that doesn't change the simple fact that there can be plenty of instances where we come to know more neutral claims *on the exclusive basis* of an inference from known less neutral claims. So, merely being a more neutral description of my sensory evidence does not entail being a more reasonable description. It all depends on what we know. It is thus simply a *non sequitur* to infer from the obvious fact that experiential claims are more neutral than external world claims, that experiential beliefs are therefore more reasonable than certain external world beliefs.⁶⁷

Can one make use of the relevant results of science to establish the point via a *reductio*? Not, of course, if it is proper to understand such results *in terms that also treat a non-cognitive theory as not obviously false*.⁶⁸ Couched in such terms, there is nothing necessarily all that epistemologically worrisome about what science tells us. According to a non-cognitive theory perceptual experience is constituted by *more* than just what happens to one from the skin on inward. Thus, the non-cognitive theorist can happily grant that science tells us that it is *nomologically consistent* for what happens to us "skinward" in perception to be held constant, while what happens externally varies so that some veridical beliefs become even radically falsidical. This is compatible with it ultimately being most reasonable to think that *any human's perceptual experience is constituted by the particular sort of nomologically consistent connection that is the (often) veridical one*. According to a non-cognitive theory, perceptual experience is indeed a process that is vulnerable to the variance that goes along with being any contingent, multiply realizable phenomenon. But, to repeat the point, if it is ultimately most

reasonable to accept a non-cognitive theory of perception, then it is just as reasonable to accept that the particular sort of contingently realized process *in our own case* is the one that is (often) veridical.

At several crucial points in the preceding remarks, I have made much of the claim that the conclusions of the above arguments do not follow given that it seems quite possible that a non-cognitive theory of perception is, as it turns out, "ultimately the most reasonable position to maintain." But, it might be asked, isn't it true that the above arguments are designed to be the instrument for establishing that very sort of reasonability? In that regard, isn't the brute *intuitive* force of the above arguments precisely why we should eventually come to reject a non-cognitive theory after proper reflection on the matter? If this were right, then it would not make sense for me to claim in *response* to those arguments that they are inadequate because a non-cognitive theory may still, in the end, be the most reasonable position. Our intuitions about the results of such arguments are precisely what are supposed to establish that sort of reasonability.

The problem, however, is that this looks at things at the wrong level of abstraction. The fact that these arguments are thought to provide strong intuitive support for a conjunctive theory reveals *strikingly* how most philosophers are already thinking of things in a way that rules out a non-cognitive theory by fiat. Thinking of things in the general way a non-cognitive theorist does, it should be clear, removes all of the force from the above arguments. For my purposes this is relevant because it shows how finding them persuasive requires, in particular, blatant violation of the totality condition as represented by (Wide+). Since, for the reasons already discussed, a properly framed first-

person, neutral standpoint requires that we be open to a non-cognitive theory's truth, we are thereby forced to admit that the above arguments do not suffice.

The above worry does, however, raise the general issue of what it would even be to philosophically defend a conjunctive theory adequately against a non-cognitive theory, if one cannot proceed with the typically offered arguments. Relatedly, another issue is how a proponent of a non-cognitive theory could show that their view is ultimately the most reasonable position, where this means, in particular, proceeding in a way that avoids the problem of question-begging while laboring under (Wide+). I will attempt to address this matter in the next chapter.

What, in any case, is the upshot of the discussion so far? One result is that we can now identify the following as the problem with both Moore's and Stroud's way of assessing the extent of our external world knowledge. In short, both rely on a claim that violates (Wide+). Moore illicitly relies on the claim that he knows he has hands; whereas Stroud illicitly relies on the claim that he does not know he is not dreaming. Stroud is right to point out that Moore violates the totality condition. But what *is* the totality condition? As I have argued, it is in part the condition that we fail to merely assume that we already know certain things about the external world. But, importantly, honoring this is compatible with *also* guarding against merely assuming that we already *do not* know anything about the external world. Obviously, when it is granted that knowledge is closed under known entailment, one way to violate this last condition would involve merely assuming that one does not know that one is not dreaming, a brain in a vat, etc. With this in mind, in this section I've tried to show why the usual arguments for the claim that we do not know such things are inadequate. At their best they involve merely assuming that a

non-cognitive theory is incorrect. Given the connection between one's views about perception and one's epistemological beliefs, *and* the impossibility involved in rebutting the claim that one is not dreaming once one accepts a conjunctive theory of veridical experience; in the context of a properly framed first-person, neutral examination of the extent of our knowledge this amounts to merely assuming that we do not know anything about the external world.

4.5 The Problem of Philosophical Skepticism

To reiterate an earlier point, the above establishes that it is at least difficult to *establish* that the sort of theory of veridical experience is correct which entails that one doesn't know that the skeptical scenarios fail to obtain. But this is different, of course, from establishing that it is philosophically most reasonable to *reject* that sort of theory. It is only by doing the latter that the problem of philosophical skepticism is effectively dealt with in favor of realism. According to this way of looking at matters, properly understood "the skeptic" is not a fallibilist. (S)he is not, that is, someone who actually sticks her/his neck out and tries to explicitly defend a conjunctive theory of veridical experience. (S)he is, rather, only someone who first tries to point out that we have no way of forcefully rebutting that sort of theory, and then proceeds to draw the appropriate negative epistemological conclusions.

This shows that the heart of the problem of philosophical skepticism lies in the mismatch between what we take our epistemic position to be, and what *second-order skepticism* entails it to be.⁶⁹ For the skeptical problem to exist, the external world epistemological realist must be unable to *establish* that it would be wrong to claim that she doesn't know that she is not dreaming, etc., and hence that she doesn't know anything

about external reality. I want to insist that this is tantamount to being unable to refute a conjunctive theory of veridical experience.

Prima facie, of course, this might make the latest version of the skeptical problem look to be just as insurmountable as it might've seemed to be in its previous form. Outright refutation in philosophy is a rare feat indeed. But I think there is good reason to be more optimistic. Though of course it would be exceedingly difficult to make good on the claim that one can catch the second-order skeptic in obvious contradiction, getting close enough to this ideal to scratch the relevant philosophical itch is, I hope to show in the next chapter, well within reach. For one, I hope it is clear that the issue is no longer centered on whether or not we can creep outside of our own skin to verify whether or not our sensory experiences "match up" with what is really the case. The thought that this sort of inability is what leads to the skeptical problem is explained by dogmatic (at least according to a properly understood first-person, neutral examination) adherence to a conjunctive theory of veridical experience. Likewise abandoned is the idea that in order to know we are not dreaming we must be able to tell "genuine" reality from "fabricated" reality in a manner similar to how we might be able to tell a real piece of fruit from a fake plastic one. This also looks at matters in a way that simply grants the truth of a conjunctive theory. It treats us as if we are *already* within an epistemological/perceptual bubble, and then asks us if we have a way of telling how things really are outside of that bubble. So, by realizing that within the context of a properly framed philosophical examination of our knowledge there is no immediate reason to accept such a "bubble theory," we realize there is consequently no compelling reason to frame the philosophical project of refuting skepticism in such fashion. This should come as a relief.

Admittedly, however, at this point the following exegetical puzzle rises to the forefront. The usual way the force of the skeptical problem is felt is as follows: it is an apparently universal view among epistemologists that there is no first-person-accessible, perceptually-based mark of the distinction between waking experience and dreaming experience. So, it is pertinent to now ask: is someone who wishes to refute the second-order skeptic still required to indicate this very sort of mark? The answer is no. Of course the non-cognitive theorist claims to know that every day or "plain" waking experience is in fact the kind of experience that involves behavioral commerce with mundane external objects. And, according to the non-cognitive theorist, what indicates *this* will be whatever it is that he appeals to in order to adequately make his case in the context of a first-person, neutral examination. So, on this view there *is* something that indicates that "plain" waking experience is in fact waking experience. But of course the non-cognitive theorist denies that sense experience itself is characterized by being trapped in a "bubble" that involves (for all we can tell) mere mental interaction with representations or other contents, and hence also that he has a way of sorting out the relevant contents into, for example, the veridical versus falsidical. He is thus in agreement with the conjunctive theorist that there is no mark of the distinction between waking and dreaming in the latter sense. But that admission poses no threat because, according to a non-cognitive theorist anyway, it rests on a false view of veridical experience.

What, then, describes the force of the skeptical problem according to this picture? It will be the fact, if it is a fact, that we are unable to convincingly show that a second-order skeptic contradicts themselves by denying that we have good reason to accept a non-cognitive theory. The issue at hand is, more accurately, properly framed in terms of

the question of whether or not human beings exemplify a particular ability to distinguish certain states of affairs from others. In the pertinent sense, to be able to distinguish fact X from some contrary fact Y means, when X obtains and it is also true that one is positioned so as to be able to come to learn that X obtains, then one is *consequently* positioned so as to be able to come to *know* that it is not the case that Y, and vice versa. Adequately motivating this view about what it means to be able to distinguish X from Y is not really necessary at the moment. What is important to realize is just that, as far as understanding the challenge of refuting a second-order skeptic is concerned, the view in question rightly sees the following to be the central task. It must not only be that a non-cognitive theory *is true*, but also that we have the epistemic tools that allow us to come to learn this, and that therefore we are able to know that it is not the case that a conjunctive theory holds. If we *did* come to know the latter in the just described way, then we would be able to *consequently* say that we can perceptually determine that our everyday waking experience was in fact waking experience, because we would know *antecedently* that a non-cognitive theory was true and hence that it was impossible for us to be dreaming, a brain in a vat, etc.

So, how on earth could we ever know that a non-cognitive theory was true? Well, of course there would be no issue at all *if* there was a *conclusive verification procedure* that we could successfully employ in a certain context in order to determine that this was the case. And here, at last, is the crucial point: if there is any verification procedure like this, then we already know what it must be. It must be nothing other than a first-person, neutral examination of our knowledge! In sum, as I see matters, in order to know that a conjunctive theory is false, and to thereby know that we can determine while involved in

everyday waking experience that we are awake and hence not dreaming, we must be able to *antecedently* determine that feeling the force of a non-cognitive theory is the inevitable outcome of a properly framed philosophical examination of our knowledge. This is tantamount, I think, to coming to know in the relevant sense that a non-cognitive theory is true. If we can't do this, then the problem of philosophical skepticism remains in full force.

Importantly, attempting to recover our knowledge of external reality on the basis of our knowledge that a certain philosophical theory of perception is true would avoid some of the usual dogmatic pitfalls of a constructive realist account. It is often pointed out, e.g., that moving from knowledge that one has hands to knowledge that an external world exists and hence that one is not dreaming, etc., looks unreasonable precisely because knowing that one has hands seems to require that one *antecedently* know one is not dreaming, etc.⁷⁰ In contrast, founding our knowledge that the mundane external world exists on our knowledge that a certain theory of perception is true engages the dreaming scenario in the proper manner. If it were true that feeling the undeniable force of a non-cognitive theory of perception was the result of a properly framed philosophical examination of our knowledge, *and* true that engaging in that sort of examination required us to deny at the outset that recovering external world knowledge was tantamount to knowing what was outside our private epistemological/perceptual bubble, then coming to accept a non-cognitive theory could not possibly amount to simply ignoring, or too quickly dismissing, the skeptical scenarios. To address the matter of second-order skepticism in the way I have in mind requires that we operate on a plane different from one that involves worrying about how to deal with the *obvious epistemic*

possibility of the skeptical scenarios. It is to think at a level of abstraction that asks, instead, whether we have to deal with the skeptical scenarios at all in order to assess the extent of our external world knowledge. We wouldn't have to if it turned out that we could establish, in the above-described sense, the incorrectness of a conjunctive theory of veridical experience. But *whether or not* we can perform that admittedly impressive feat is precisely what we would be out to discover in performing our "external" examination.

This means that in order to properly assess the prospects of a philosophical defense of epistemological realism about external reality, we have to get extremely clear on just how a first-person, neutral examination of our knowledge must proceed, especially when we are laboring under (Wide+). In this chapter I have been helpfully vague at times in talking about the positive aspects of such an examination; though given my primary aim I have said quite a lot about how it cannot proceed. In any case, it is clear that for the realist to be successful the examination in question would have to allow one to both fairly engage a second-order skeptic, as well as forcefully motivate a non-cognitive theory. In the next chapter, I will state what I think is the realist's best shot: an appeal in a certain carefully described theoretically-neutral way to the *presentational* or *transparent* phenomenology of "plain" waking experience.⁷¹ When the accepted obviousness of a conjunctive theory of veridical experience is no longer weighing one down, I think such phenomenological appeals can provide significant positive, non-question-begging support for a non-cognitive theory. Let me now try to make the case.

CHAPTER V SECOND-ORDER SKEPTICISM REFUTED

5.1 Introduction

Given the result of the previous chapter, we can say that the epistemological realist has this much to say for her view: if the above is correct: *for all you and I know* the skeptical scenarios are mere flights of fancy. At this point in the examination, after all, for all we know a theory of perception which identifies these scenarios as epistemically impossible is true. Nevertheless, this result is quite compatible with second-order skepticism. The openness of the last "for all we know" remark reveals that a conjunctive view of veridical experience is also at least still on the table. At this point we can say that it is equally true that for all you and I know the skeptical scenarios *are* epistemically possible. Thus, rather than showing that the problem of philosophical skepticism is resolved in the realist's favor, the arguments above even seem to clarify and hence reinforce it. Someone might quite rightly point out that it is precisely *because* we can't say, with good reason, that a conjunctive theory of veridical experience is false which explains why the problem of philosophical skepticism exists.

5.2 A Non-Cognitive Theory of Perception

The relevant question therefore becomes: can one avoid the realist's dilemma when merely laboring under (Wide+)? Can one, that is, provide a non-question begging argument for a non-cognitive theory of perception given what (Wide+) dictates regarding what does/does not constitute question-begging? But there are two different projects that must be distinguished here. There is the project of avoiding the realist's dilemma when *illustrating* the methodological possibility of a non-cognitive view; and then there is the project of avoiding it when *arguing* for the claim that this sort of view is the one that we

ought, as philosophers, to maintain, rather than the other available methodologically possible views. I will spend the rest of this section showing the former. I will show, more particularly, that it is not a *methodological necessity*--when performing the "distinctively philosophical" sort of assessment that the skeptic carries out--that one adopt a form of a conjunctive theory out of a need to avoid the realist's dilemma. It should be noted, however, that illustrating this is compatible with what I think is an obvious feature of the assessment I am about to offer; namely, its being a kind of "Just So Story." The goal here is only to show the methodological possibility of the tale I am about to tell. The account offered obviously retains this characteristic even if it also looks *ad hoc* when placed within the context of this chapter. Rest assured that arguments for this view--one's designed to also steer between the horns of the realist's dilemma--will be given in the next section.

So, suppose that at a certain time t_1 I honor (Wide+) and fail to assume that I have any knowledge of the external world. At a later time t_2 , however, suppose I come to accept, for reasons I will discuss, the following as a position I find not only plausible, but quite relevant for the sort of assessment at hand: *prima facie, one primary function of my sense-organs is to allow me to use my body to interact in various ways with an independently existent world.* Unless one is already a card-carrying conjunctive theorist, coming to accept this view can be seen as one natural outcome for the sort of assessment in question. One could first recognize as Descartes did (initially) that any knowledge of the external world that we might have must be knowledge gained *via* our sensory experiences. One would then be inclined, of course, to search for any features of one's sense-experiences themselves that might be relevant. At this point if one had the

pervasive conjunctive theorist bias, one might (as Descartes did in his own way) come to think that the relevant feature is that of *providing* one with *information about* external reality. This is indeed a way to proceed, and hence we must acknowledge the methodological possibility, when laboring under (Wide+), of a conjunctive theory of veridical experience. Nevertheless, I am suggesting anyway, one could just as consistently see things as I have above. If one were, like I am, antecedently rather suspicious of the metaphorical flavor that conjunctive theorist accounts of sense-experience have, then all I am suggesting here is that one could instead accept the view that one's knowledge must be a *function* of what one's sensations *do* for one, in a manner that pertains in no direct way to what sorts of *beliefs* it furnishes one with and/or *constitutively* underwrites. Note, of course, that this is a conception that is wholly alien to the conjunctive theorist's way of seeing matters. Sensory experience, on this view, is not about providing us with information. Instead, it is this or that function of one's sensory experiences with respect to *guiding one's bodily behavior* that is thought most relevant, on this account, for addressing the question of one's interaction with, and hence the extent of one's knowledge of, external reality. But of course, examining the function of one's senses leads naturally to examining what it is that performs the function. This could then lead one to examining the role one's sense-organs themselves might play in providing one with any knowledge of external reality.

In the account just given I have come to explicitly countenance, among other things, my sense-organs as things that exist. It is worthwhile to pause here and allay any worries that I have broken the rules. *Coming* to accept the above view nowhere involve my "merely assuming" that I have sense-organs, or anything else that might help to

explain (down the road, say) why I have a lot of external world knowledge. We must remember that what is and is not a case of "merely assuming" is a function of the shared background assumptions of the disputants involved. Note, then, that the charge that one, in coming to this assessment, has "merely assumed" one has sense-organs would, if thought to be an effective polemical point, allow for a *parallel* charge to be leveled against any conjunctive theorist who, for this same sort of assessment, came to accept the view that their sensory experiences primarily function so as to provide them with information about external reality. I say this only because it illustrates that *both* charges, under (Wide+), are out of place. Coming to accept either the above-described view *or* the conjunctive theory of veridical experience would not, in either case, allow one to be subject to the charge of "merely assuming" anything. This is because the methodological perspective at work here is one *neutral* between the question of whether it is the non-cognitive theorist's or the conjunctive theorist's "rabbit" that is "in the hat already."

So, to resume the thread of the "Just So" yarn I have been spinning, it is of course easy to see how one can get from the view that it is *some* function of one's sense-organs that is relevant for addressing the question of one's knowledge of external reality, to the plausible enough view that one's sense-organs have as one of their apparent *primary functions* that of allowing one to comport oneself appropriately in an independently existing world. Perhaps before coming to accept this view, one noticed that if one was moving around but not looking at where one was going, one would often bump into things; whereas if one's eyes were open, those same sorts of things could often be easily evaded. Perhaps one noticed how hearing a buzzing sound grow in an appropriately intense crescendo allowed one to *track* when an object that was the apparent source of

that buzz was getting nearer. Such examples could be multiplied indefinitely. The basic feature of one's sense-experience that this sort of account would be hammering at is clearly *whatever it is* about one's sensory experience that allows one to evade and/or come into contact with environing objects of the appropriate size, shape, speed, etc. Suppose, then, that I come to accept such a view at some point in between times t_1 and t_2 .

So much, I think, for illustrating the methodological possibility of a non-cognitive theory of perception. It is easy to see, of course, that *if* the arguments of the next section are successful: if the view that my sensory experiences *fulfill* this function (i.e., actually do allow me to interact with an external world) were somehow shown to be the view one ought to maintain, then this would allow me to say I can deduce the falsehood of the skeptical scenarios. If I ought to maintain that a primary function that my sense-organs (and hence sense-experiences) *continuously fulfill* is to allow my body to interact, consciously or "unconsciously", with an external world, then this implies that it is just as reasonable to suppose there is an external world that my body (including my sense-organs) interacts with, and hence that any view which implies that there *might* not be ought to be considered false. This would leave many questions about the extent of our external world knowledge unanswered, of course; some of which I will take a stab at answering in the next chapter. But it would nonetheless rebut *fallibilism*, and hence resolve the problem of philosophical skepticism in the epistemological realist's favor.

5.3 The Main Argument

It is thus towards the all-important matter of overcoming second-order skepticism that I will now turn. As a way of attacking the issue that will aid my own purposes, take

the following version of a familiar way to describe the *paradox* of skepticism. The three claims:

- (1) I know that my sense-organs allow me to interact with an external world.
- (2) I do not know that I am not merely dreaming that my sense-organs allow me to interact with an external world.
- (3) I know that my sense-organs allow me to interact with an external world only if I know that I am not merely dreaming that my sense-organs allow me to interact...

all have their fair share of plausibility; and yet, of course, they all can't be true. This is simply an abbreviated, extremely simplified way of stating what I have called the problem of philosophical skepticism. Henceforth, let's call this *the skeptical paradox*. The non-cognitive theorist should obviously try to offer support for (1) (and (3) for that matter) if possible, in order to defend denying (2). The second-order skeptic will maintain, however, that no such support for (1) can be forthcoming. Obviously, then, the matter hinges on whether the anti-skeptic can offer adequate support for (1) (and, as we shall see, certain other knowledge claims). I will now try to do so in two stages. In the first stage, I will present an argument designed to convince anyone who, unlike the second-order skeptic, agrees that the skeptical paradox exists; that rationality demands that we see a way out of the relevant difficulty. I then, in the second stage of my overall argument, make a case for the claim that the argument of the first stage does not, despite initial appearances, misfire as a response to the radical skeptic.

Given the presuppositions about who it is designed to convince, it is no surprise that the key premise in the argument of the first stage is simply: *there is something prima facie paradoxical about the skeptical result*. In combination with earlier results, this premise lends substantive support to a non-cognitive theory. If one thinks the skeptical paradox exists, then this can only be because, along with (2) and (3), one thinks that

denying (1) would be absurd given its prima facie plausibility. But *if* one's sound reasons for denying (1) depend only on what one takes to be sound reasons for *affirming* (2), then even granting the latter claim's prima facie plausibility, it follows that the arguments of the previous chapter shows that there is no good reason for denying (1). (That is, just so long as those arguments do provide grounds that sufficiently weigh against (2)'s *prima facie* truth). In other words, since the existence of a paradox requires there to be an unhappy relationship between the above three prima facie plausible claims, once we have found as a result of the previous chapter that there is no convincing case for (2), we can use this as a reason for saying there is thus no convincing case for denying (1). This, plus the additional premise that denying (1) is prima facie absurd, allows us to conclude that--barring the subsequent introduction of other countervailing reasons for denying (1)--we ought to simply affirm (1). This result could in turn provide the basis, for its intended audience anyway, for rejecting (2) as simply *false*, along with other similarly structured claims that make (on this view, unwarranted) reference to the epistemic possibility of systematic deception. Of course, it would be somewhat misplaced to say that the above arguments show we now *know definitively* that (2) is false (whatever that could even mean). Instead, the matter is most clearly described as follows: according to the view of one's sensory evidence we have determined (via the above argument) we ought to maintain, (2) is false. I will call this the *argument from paradox* for a non-cognitive theory of perception (i.e., the truth of the claim found within (1)), and hence for the claim that we ought to think we know we are not dreaming, a brain in a vat, etc.; in short, for why we ought not be fallibilists.

The argument from paradox has, I admit, a spurious ring to it. Nevertheless, I do think it contains the essence of the proper answer to the problem of philosophical skepticism. Since I anticipate the citation of the crude similarity it bears with such arguments, I am somewhat called upon to address its alleged "Moorean" flavor. I will do so at the end of the chapter. But before doing so I need to examine, of course, how the second-order skeptic in particular will find fault with the argument from paradox, and thus begin the next stage of my overall argument. The second-order skeptic's complaint will be that, on his view, there is simply no paradox in the skeptical result, because it is not clear, he thinks, (aside from, perhaps, (3)) why we ought to affirm any of the relevant claims. Of course, he can recognize the *relative* tension between the above three claims. But he will fail to see how that tension might be utilized to serve the purpose I require of it. He will see it as, at best, a form of "boot-strapping" one's way to the reasonableness of a certain position. The skeptic, being a doggedly neutral party, will maintain a view that is defined by failing to be able to see even the potential that any argument from paradox might have for supporting a certain theory of perception, and hence for deciding, in part, the matter of the extent of our knowledge of external reality.

For this reason I will spend the rest of the chapter making the strongest case I can for the view that the skeptic can't *consistently* deny the absurdity of denying the claim found within (1) in particular--that my sense organs allow me to behaviorally interact with an external world--(as well as claims found within other knowledge claims); not, that is, without abandoning his powerful argumentative position. Since, as it turns out, the reasons for finding denying claim found within (1) to be absurd are *also*, when combined with the non-cognitive theorist's argument from paradox, the reasons for not

being a second-order skeptic, this response, if sound, will show why the skeptic cannot remain a skeptic while also avoiding succumbing to the argument from paradox. If I am right, then as it turns out we can't think that the problem of philosophical skepticism still exists after we have properly examined the matter. We "can't" that is, just to the extent that I am right in thinking that the second-order skeptic can't consistently deny the absurdity of denying the claim found within (1).

To see why the second-order skeptic, in particular, has to admit that denying the claim found within (1) would be absurd, we are required to identify and scrutinize what one might call the "normative source" of that absurdity. The idea is to explain why it is plausible to claim that even the skeptic must recognize the presence of this source given the neutral methodological perspective currently assumed; or otherwise, if we are thought unable to accomplish this, indicate what position he is thus left in.

The first thing to realize is that clearly the source of the absurdity in denying the claim found within (1), as well as those found within similar knowledge claims, is *phenomenological*. This is made evident by noting the following contrast in the treatment of *knowledge* claims: those moved by the cognitive dissonance of the skeptical result are much more likely to find puzzling the denial of claims like "I know I have hands" as opposed to those like "I know she'll be at the party." In the appropriate sort of context, denying the latter, in a way quite unlike denying the former, would *not* be thought to be absurd. People simply play fast and loose with the term "know" all of the time. But failing to know in many of the fast and loose cases is compatible with the denial of the claim that I can't come to know plenty about the external world after a more careful examination. Even if I know that I don't know she'll be at the party, I may, contra

what the second-order skeptic claims, still know that I know an external world exists and hence that I am not dreaming, etc. It is somewhat obvious, then, that the particular kind of claim one might find within a certain knowledge claim that is relevant for giving rise to the skeptical paradox, is one that has an intimate relation to our sensory experiences themselves. More accurately, the claims that really matter seem to be of the sort where *the only intelligible way* their truth can be questioned (in an otherwise typical seeming circumstance) seems, to anyone who scrutinizes the matter, to be by citing the supposed epistemically possibility of skeptical scenarios describing *systematic deception*. It is for this reason that denying *knowledge* claims like (1) and "I know I have hands" qualifies as paradox-engendering. In contrast, a claim like "I know that she'll be at the party" is one that can be shown false, albeit in a rather nit-picky-sounding manner, by citing epistemic possibilities much more palatable and every day. Indeed, it warrants repeating that admitting even the ubiquity of epistemic possibilities like the latter is compatible with the falsehood of *second-order skepticism*.

Obviously, then, the difference is that the paradox-engendering knowledge claims are ones that contain claims of a sort (primarily) prompted by and (primarily) maintained as a result of our on-going sensory experiences. I do not think there is a precise line that exists between those knowledge claims that are prompted and maintained in this way and those that are not. Nevertheless, it is equally clear that no one can feel the sort of cognitive dissonance that the skeptical paradox engenders without having the denial of the sorts of knowledge claims discussed above in mind. The hallmark of those claims, I am suggesting, is the large enough degree to which the claims they contain are prompted and maintained by one's (often non-linguistic) sense-experiences.

As a result, this invites the question: what feature(s) in particular of our sensory experience can be said to give rise to the absurdity of denying knowledge claims like (1), that I know I have hands, and the like? Well, since it is of obvious relevance, we should here note how the non-cognitive theorist would phrase her *description* of such features. Denying (1), etc., would be absurd, according to this view, as a result of the feature(s) of our sensory experience which make(s) us poised to comport our bodies in various ways in our surrounding environment. Of course, this should be taken in a way that does not read any potential non-cognitive theorist as claiming that her sensory experiences *represent* us as so interacting with our environment. The whole point of citing the methodological possibility of a non-cognitive view of perception in the first place is to avoid both what one might call *the perceptual content wars*--i.e., the debate that takes it as granted that our perceptual experience has content, and is thus over what kind of content our perceptual experience has and whether that content allows us to certify the truth of our external world beliefs--and some of the consequent reactions that presuppose its terms. I am suggesting that it is, in particular, by way of first recognizing the methodological possibility of a non-cognitive position that one can come to achieve a perspective regarding the matter of sense-experience itself which, if plausible, would imply the implausibility of their being anything like The Realm of Private Experience, The Cartesian Theatre, or any other member of that lot of metaphors (presiding over all of which, I think, is the idea that perceptual experience has "content"). Shed of such skin, the *non-cognitive theorist's* construal of the absurdity of denying that I know a fast-moving object that looks like a bus is heading directly my way (in an otherwise typical seeming circumstance), e.g., is that such absurdity stems ultimately from recognizing (in

the context of the skeptic's favored sort of assessment) the non-linguistic features of my sensory experiences themselves which make it so that I am poised to move out of a way in time so as to avoid being hit by it. I am poised for that sort of movement, on this construal, only because I am aware in some sense of what my own relation is with other objects in my immediate environment that are of the appropriate size, shape, speed, etc. Of course "aware" is a term that has a cognitive use, and thus one should not be misled at this point. I mean the term "aware" here in a way that is not supposed to reference any particular *cognitive, belief-producing* or *belief-justifying* function of my sensory experiences. I think anyone will be able to see the feature I am speaking of in this way if they are also one who has *not* already engaged and chosen a side in the perceptual content wars (by simply granting that perceptual experience has, say, representational content and hence choosing, say, either representational realism, skepticism, or phenomenalism for that matter). The presence of the features of our sensory experience that allow us to move our bodies out of the way of passing buses (among other things); *that*, the non-cognitive theorist can insist, explains why it would be absurd to deny (1) and similar knowledge claims.

It follows, then, that plausibly denying that denying (1), etc., is absurd, as the second-order skeptic clearly must, ultimately requires him to either deny the presence of the relevant, paradox-engendering non-linguistic features of our sense-experiences, or deny that we have an epistemic right to privilege the non-cognitive theorist's *description* of those obvious features. Of course, if the skeptic pursues the first option, he is in trouble. The non-linguistic features of human sense-experience that the non-cognitive theorist speaks of in a certain way so as to make the absurdity of denying (1), etc., clear,

are simply present in human beings, and their presence is able to be verified by anyone sophisticated enough to do so, and who (perhaps) is also familiar enough with the relevant ways of speaking and hence distinguishing certain things from others. Denying the presence of such features would require that the skeptic fail to be honest about the nature of his own subjective experience. Thankfully, in any case, "the skeptic" as often described does notice the features in question. The skeptic of Descartes' second *Meditation*, e.g., certainly did. *That* sort of skeptic would note well the non-linguistic phenomenological features I have been citing, he just would choose to *phrase* his description of them in a conjunctive-theorist-friendly fashion that, when adhered to consistently and systematically, removes all inclination to believe that the relevant sort of absurdity exists. This sort of skeptic would say something like that at the relevant time he was having, e.g., sensory experiences *as if* a bus was coming at him, etc. What is noteworthy here, however, is that what follows the "as if" in any such proposed description, since it is *about* the external world, shows clearly that this sort of skeptic recognizes, albeit in his own way, the presence of the non-linguistic phenomenological features in question.

This brings us, of course, to the second way the skeptic can respond; he can simply reject the claim that we have good reason to privilege the non-cognitive theorist's *description* of the non-linguistic features in question, over a contrary conjunctive theorist's description. He can provide the following reasons for this: under the conjunctive theorist's description there is no obvious absurdity in denying (1)--given, of course, the intimacy of the logical connection between that theory of perception and reasonably affirming a claim like (2)--and, furthermore, we still seem *free*, within the

constraints of our chosen methodology, to adhere to this latter sort of description. It thus follows, according to the skeptic, that we have not yet been presented with a good enough reason for denying (2). Nothing about the distinctively philosophical assessment *forces* us to deny (2).

Now, quite obviously, the skeptic is free to reject a non-cognitive theorist's description for the above reasons. What is important to see, however, is in what shape this leaves his position. Anyone who has played by the skeptic's rules up to this point, and who recognized *from the neutral perspective so adopted* the non-linguistic features of their own sense-experience which, when described in non-cognitive-theorist-friendly fashion, make the absurdity of denying (1) and similar knowledge claims clear, can despite all this, still admit that one is in general *free* to choose to instead phrase the relevant description of such features in the conjunctive theorist's favored terms. And consequently, this philosopher must in turn acknowledge that if one remains consistent and systematic about phrasing the relevant features in conjunctive theorist terms, one can remove or disguise the presence of the relevant absurdity. But as Wittgenstein was quite adept at pointing out, the relationship between experience and language doesn't work so that experience itself *forces* certain descriptions on us. And thus, any philosopher who has played by the skeptic's rules in the above sense and is *consequently* a non-cognitive theorist can admit that, nevertheless, *after* anyone else got accustomed to speaking in a conjunctive theorist's favored fashion, it could even be that the latter person no longer thought that the phenomenological features in question bespoke the absurdity of denying (1) and similar claims. But such an attitude is clearly one that is taken only after one has gotten accustomed to the conjunctive theorist's way of talking. The non-cognitive

theorist can rightly point out that there is a kind of re-conditioning involved here; *a kind of re-conditioning that does not occur at all in the case of coming to adopt the non-cognitive theorist's description in the first place, for one involved in the relevant, neutral philosophical examination.* The non-cognitive description, unlike that of its counterpart, is a much more *natural* way of describing the basic “transparent” phenomenological features I have been harping on. In contrast, we must clearly learn to speak of things as the conjunctive theorist prefers, and only as a result of so learning come to accept that there is nothing absurd in denying (1). But as long as one was able, before being re-conditioned, as it were, to recognize the non-linguistic features in question as lying at the source of the absurdity of denying (1), then this re-conditioning obviously points to a way in which anyone who chose to always make use of their right to use conjunctive theorist language would only be, in a sense, fooling themselves about the “rationality” of their subsequently determined view with respect to denying (1) and similar claims.

The important feature of this last response is how weak it makes the second-order skeptic look subsequently, since obviously he is someone who would want to treat the openness in how to describe the non-linguistic phenomenological feature I have been harping on as the crux of a very important epistemological matter. He would claim, I guess, to be unable to see why we can say that someone who has now come to adopt a non-cognitive position *is free* to respond to the admissibility of a conjunctive theorist's description in the above, admittedly flippant-sounding way. The skeptic, being unable to adopt the sort of perspective that allows him to see how such epistemological wheels turn, will persist in the thought that all of the above is idle. And alas, I admit that it is difficult to know what to say here. But, nor does my failure to know how to continue

discourse bespeak dogmatism. We have, after all, the following reasons for adopting a non-cognitive theory of perception and hence for claiming we can resolve the problem of philosophical skepticism in favor of epistemological realism: (i) a non-cognitive description is a methodologically possible account of our sensory experience, in the sense given in the previous section; (ii) such an account avoids engaging in the *quite unnecessary* perceptual content wars; (iii) if we adopt it we are able to avoid the absurdity of denying (1) and similar, paradox-engendering knowledge claims; and (iv) no independent, countervailing reasons seem to exist for denying such claims.

Substantively, then, the whole matter *may* come down to this: (iii) of the previous paragraph can be thought to beg the question as a means for supporting a non-cognitive theory, but only if it is implausible to deny that *the only way to see the absurdity of denying (1) is to have already accepted a non-cognitive theory of perception*. But, thankfully, this last claim is not only able to be plausibly denied, it is even quite clearly false. The matter pertains to what I have already touched upon above; regarding which description of the relevant phenomenological features involves a re-conditioning about one's attitude with respect to the absurdity of denying (1) and similar knowledge claims, and which does not. Notice, after all, the following rather important asymmetry. It is quite obvious that there are *no* non-linguistic phenomenological features that, by themselves, one can cite as providing adequate grounds for thinking, at any given time, that one *might* be dreaming, a brain in a vat, etc. Somewhat ironically, this can be supported as follows: even while one is dreaming, one (usually) does not think that one *might* be dreaming, on the mere basis of what is happening to one at that time. To put the point in another way, the claim that *I know that my current sensory evidence does not*

guarantee that I am not dreaming is clearly *not* one that is a member of the above-described, paradox-engendering lot: it is *not* one, in other words, that contains a claim that is sufficiently prompted and maintained by one's ongoing sense-experiences. Philosophers seem to forget that anyone who thinks otherwise has to *both* be someone with normal sensory experience as well as someone who has heard some fanciful-sounding stories while in a certain mood. In contrast, as I've already pointed out, there is no similar peculiarity in claiming that one's sensory experience does, by itself, seem to give one adequate reason for thinking that the truth-conditions for certain beliefs about the external world are simply met, and thus have no way of being false, and thus are known. It is not like I need to be told a fanciful-sounding and elaborate *story* to, while engaged in a typical sensory experience, *also* find it at least *prima facie* plausible to think that it is *simply true* that I have hands, that I am walking around in a room, etc. And thus, it certainly follows as well that I do not need to be first indoctrinated with a certain complex philosophical theory of perception (non-cognitive or otherwise) in order to find such claims obvious. Instead I need only to be someone with a typical kind of sensory experience.

Of course, I do not expect the above to convince anyone with antecedent conjunctive theorist sympathies. Thus, I do not expect it to convince the admittedly large lot of philosophers who seem to treat it as obvious that subjunctive conditionals like the following hold: *Were I to be dreaming right now that p, then from a purely subjective standpoint I wouldn't know the difference.* These philosophers need to consult the previous chapter to see why I find their position extremely dubious. I do hope, in any case, that the argument can be found convincing to those who do not think it is obvious

that they don't know they are not dreaming, etc., and who are still wondering whether or not they ought to be second-order skeptics.

The skeptic, being "doggedly neutral" will, however, still not be moved by the above pleas. Thus, the one open to the matter of second-order skepticism still will have skepticism as an option. But anyone who chooses skepticism at this point has, at least in my view, carved him/herself out a very remote corner in logical space. Note that he/she is akin, we may assume, to the person moved by the above arguments in the following sense: we can say that he has first adopted, while considering the matter, a neutral stance with respect to the tenability of either a non-cognitive theory or its conjunctive counterpart. But he is unlike the person subsequently moved by the above argument in being unable to also "see" the absurdity of denying (1) *from that neutral stance*, when the presence is noted of certain non-linguistic phenomenological features of his/her own experience. Surely, though, the *bare* claim that one "just can't see" this is a very unconvincing reason to remain neutral (right?). In contrast, the admittedly parallel ability I require of "just seeing" or "recognizing" that it would be absurd to deny (1), has the following to back it up: namely, the primitive non-linguistic phenomenological features of one's own sense experience that, I have argued, can be "recognized" as explaining the absurdity of denying (1) from a neutral methodological perspective. The point, of course, is that it follows that the methodology involved plays by the skeptic's rules, since, after all, no questions will have obviously been begged. It thus follows that the non-cognitive theorist's position so arrived at seems, anyway, to be as far from dogmatic as one of its kind could get.

5.4 Conclusion

In sum, the above shows you cannot adopt a second-order skeptical position except when it is one of the following extremely anemic form. It is a position that says we ought to still, despite all of the above, remain neutral on the matter of how to phrase the non-linguistic feature of our sensory experience that, when phrased in the non-cognitive theorist's way, makes quite clear the absurdity of denying (1) and similar knowledge claims. This position is a far cry, indeed, from the powerful sort of skepticism that Descartes tried to motivate in his first Meditation.

Is my position just an elaborate version of Moore's famous way with skepticism, and thus subject to the same problems? Doesn't the central role of the crucial premise in the argument from paradox show that I am simply *assuming*, albeit in a primitive and (I claim) neutrally-motivated way *that I have the knowledge that an external world exists with which my body continuously interacts?* Well, we must yet again recall that talk of what is and is not question-begging makes sense only when set against a background of certain shared assumptions between the disputants. The key move for me, of course, is the claim that one can notice, *from a standpoint that treats whether it is a non-cognitive theory or a conjunctive theory that is correct as an open question*, the phenomenological features of one's own experience which themselves allow one to subsequently "see" that it would be absurd to deny (1) and similar knowledge claims. Moore did not look at the matter in precisely this way, to the best of my knowledge. It is true, however, that one way of (tendentiously) describing the matter is that I "start" with a *fundamental recognition*--the absurdity of denying (1)--which is of a "Moorean" nature. But when this "fundamental recognition" is placed within the context of the philosophical examination

of our knowledge as I have placed it, it is completely false to say that I am *simply* starting with that view as an unargued for assumption, and then going on from there to refute skepticism. If there is any "starting place" of any sort in my overall case against the skeptic, it is with the basic *non-linguistic* phenomenological features whose mere presence, I have maintained, are the source of the absurdity of denying (1) and similar knowledge claims. Where Moore and I differ, then, is not in terms of the sort of knowledge claims we hold true which thereby allow us to say what we want to say about the skeptical scenarios, but rather, it seems, in terms of the particular manner that we would motivate the epistemological worth of these claims, and also in terms of what role we would have such claims play in the overall case made.

CHAPTER VI INFALLIBILISM DEFENDED

6.1 Introduction

Having defended my central position, the primary aim of this concluding chapter is to provide an important development of it: a discussion of how it is that we come to know certain things about the external world with complete evidential certainty. Since I see the arguments given above as able to stand on their own feet, the remarks that follow will sometimes be provisional, and in certain cases exposed explanatory gaps will remain unfilled, providing merely suggestions for future work. Before providing this somewhat partial exegesis of the connection between a non-cognitive theory and infallibilism, however, for the purposes of better connecting up my position with contemporary views I will first explain what perspective I think the *assumed truth* of it gains on certain thought experiments which have played a central and divisive role in contemporary epistemology; specifically, the epistemological internalist/externalist controversy. I will examine both the New Evil Demon thought experiment (which I will assume is designed to confirm internalism and disconfirm externalism) as well as the case of forgotten evidence (which I will assume is designed to confirm externalism and disconfirm internalism). The goal here will be to suggest that--quite unlike popular versions of both internalism and externalism--the intuitive force of *both* of these thought experiments can easily be understood to be consistent with an infallibilist picture of what our epistemic position with respect to external world claims can sometimes be. I should stress here the somewhat obvious point that the goal will not be to show why these thought experiments *confirm* this picture. The confirmation of infallibilism resides, rather, in the arguments

given in the previous chapter, as well as the (partial) explanation offered in the final section of this chapter.

6.2 Fallibilism, Internalism and Externalism

Contemporary epistemological internalists and externalists with respect to epistemic justification are, as far as I can tell, uniformly fallibilists. Fallibilism lies at the foundation of their whole dispute as a shared presupposition. By questioning the grounds for fallibilism, I have thereby questioned the very sense in thinking we *must* go one way or the other with respect to that controversy, at least as it plays out in the contemporary literature. Both internalism and externalism are, as contemporarily stated, without sufficient support because both rest on fallibilism. However, if this correct, then we are left with the following issue. Denying that we must make a choice between these views leaves a kind of meta-epistemological lacunae that must, it seems, be filled. How should we *plausibly* understand the matter of the justification of external world beliefs, if not in the currently accepted fashion?

Here's what I propose: we must maintain, as a fundamental methodological principle, that a suitably general understanding of *perceptual* experience requires us to focus only on *veridical* experiences--experiences where spatio-temporal objects play the appropriate kind of regulative role. And yet, stating this response only seems to draw the resultant meta-epistemological lacunae into greater relief. There are *extremely* well-entrenched conceptual reasons, after all, for calling systematically hallucinatory experiences a kind of perceptual experience; and hence for disputing the very sense in attacking a general understanding of perceptual experience in the above-described way. Of course, if what I have argued above is correct then these conceptual reasons would

depend, for their forcefulness, on persisting in an idiom friendly *only* to a conjunctive theory of veridical experience. We have seen, that is, that it is only a conjunctive theory of veridical experience which could deliver the result that some perceptual experiences (namely, the ones involving systematic deception) aren't veridical experiences. An infallibilist, in my view, should simply not speak in that sort of way about perceptual experience. But then the issue is this. If we ought to reject that way of speaking, what should we make of the forcefulness of the traditional thought experiments employed to illustrate the relevant conceptual points? If the infallibilist had *nothing* to say here, then this would provide powerful reasons for persisting in an idiom friendly only to a conjunctive theory, and hence for holding that some perceptual experiences are hallucinatory experiences. Furthermore, and much more troublesomely, for the infallibilist to simply *dismiss* that idiom is, it seems, *just another form of dogmatism*. In fact, it is a round-about, albeit polemically rather clever, way of essentially *agreeing* that the problem of philosophical skepticism exists. At a general enough level of description, after all, the skeptical problem can be said to dawn on one once one realizes that the *best one can do is be dismissive* when engaged in argument with someone who maintains a certain putatively consistent position that is also contrary to "common sense."

Broadly understood, there are two kinds of cases worth discussing. There are cases that are designed to support internalism to the detriment of externalism, and vice versa. I will now examine one powerful instance of each of these cases, in order to show that one can, in fact, talk plausibly enough about them in a way does not require adopting an idiom friendly only to a conjunctive theory of veridical experience. This will, in effect, back-fill the meta-epistemological lacunae left by rejecting fallibilism in a way

that takes seriously enough the intuitions about justification apparently confirmed by these thought experiments.

A powerful instance of the kind of case designed to confirm internalism to the detriment of externalism is the New Evil Demon thought experiment.⁷² This thought experiment asks us to imagine two subjects who are exactly alike phenomenologically speaking. What it is like to be the one subject is type-identified with what it is like to be the other. It is just *also* that one subject is having a veridical experience while the other a vivid hallucination induced by systematic deception. We are then asked whether these two subjects have an identical kind of evidence for a certain perceptually-prompted/maintained belief. The powerful intuition that one must either (as the internalist does) honor or (as the externalist does) attempt to explain away is that they *do* have the same sort of evidence for that kind of claim. But if one accepts this, it would follow that being in what the non-cognitive theorist describes as the relevant sort of relationship with spatio-temporal objects is not required in order for one to have phenomenologically-based evidence *of some kind* for their presence.

In suggesting that we need not make a choice between internalism and externalism as typically espoused, I am, it seems, required to deny that there is a perfectly legitimate question posed by this thought experiment. But the thought experiment certainly *sounds* coherent, and hence the question of whether or not their perceptually-based evidence is of the same kind certainly *sounds* legitimate. And, of course, *if* the question is legitimate, then contrary to what I have claimed, one must choose between looking at evidence either internalistically or externalistically; between answering yes or no. The fact that there clearly seems to be a legitimate question asked here, therefore,

provides a powerful conceptual reason for engaging in the fallibilist/conjunctive theory of veridical experience line of inquiry.

Obviously, though, right at the outset an infallibilist *can* legitimately question the above thought-experiment's very coherence. She could explain the prevalent thought that it asks a legitimate question as merely betraying an arbitrary adherence to the fallibilist/conjunctive theory line of thought. There is simply no compelling reason to think that we can have two subjects who are phenomenologically-speaking identical and yet non-identical (in the relevant sense) in terms of what happens "outside their skin;" unless, of course, one is *already* thinking in a way friendly only to a conjunctive theory of veridical experience.

Equally obvious, however, is that a more theoretically neutral version of the thought experiment can be easily formulated. The infallibilist would have to, it seems, accept that there can be two subjects who are having experiences that are "internally" very much alike (i.e., in terms of what is happening to each of them inside their skin); but where one subject's experience is the result of the proper regulative role of actual external objects, while the other's is the result of merely having his sensory organs stimulated in some deviant way. Filling the meta-epistemological lacunae thus requires the infallibilist to show that any worthwhile meta-epistemological lessons that can be learned from *this* formulation of the thought experiment are consistent with denying fallibilism.

Thankfully, nothing requires one to say that the two subjects are having identical *phenomenological* experiences in this kind of case. They are having experiences that are "internally alike;" that much is true. But if in phenomenology we are after a suitably general understanding of *perceptual* experience; and if in that sort of study we need *not*

account for hallucinatory experiences, then among other things we need not worry about the fact that such experiences are internally similar to veridical experiences. The similarities here are obvious, and hence must be acknowledged; but this is compatible with claiming that they are *irrelevant*. Analogously, one could acknowledge the similarity that any painting bears to any *blank* canvas, but nonetheless deem that a completely blank canvas could never serve as a counter-example to a proposal for a general understanding of *painted* art. Any blank canvas would, let us suppose, have *some* of the necessary conditions for being painted art; but it would obviously lack others. Similarly, a vivid hallucination can be granted as meeting *some* of the conditions for being a perceptual experience (namely, those that pertain to what happens from the skin inward) but as lacking others. Of course, the fact that it is compatible to admit that there can be two subjects who are involved in the cases described in the neutral formulation of the thought experiment, but deny that they are identical *phenomenologically speaking*, shows that the thought experiment by itself fails to confirm fallibilism, or any view about perception friendly only to it.

Nor does admitting that we can *conceive* of such hallucinatory experiences alongside internally alike veridical cases, by itself, show that human perceptual experience is not constituted, in part, by the proper regulative role of external objects in one's surrounding environment. Admitting this kind of conceivability, in other words, is quite compatible with thinking that *what it is like* to be a normal *perceiver* is not conceptually separable from what it is like to track, to a degree of varying success and in various ways, spatio-temporal objects in one's proximal external surroundings.

Nor, finally, does the intelligibility of this neutral formulation of the thought experiment, by itself, suggest that the subject in the veridical case "can't tell" that he isn't in an appropriately understood hallucinatory case. If by merely knowing he is having a normal, everyday "plain" perceptual experience the subject knows that he is having a veridical experience (as a non-cognitive theorist would define this term) then by merely being a conscious subject one could tell one was not the victim of systematic deception. In sum, far from forcing us to choose, when the New Evil Demon thought experiment is formulated in neutral terminology it becomes clear that, by itself, it provides no compelling reasons for accepting the claim that the two subjects are having identical *phenomenological* experiences. If that is true, then for the reasons discussed above it provides no reason for thinking we have to choose between internalism and externalism.

But perhaps the infallibilist can be thwarted in a different way. What is it about justification that *is* confirmed by the neutrally formulated thought experiment, if it is not something that is friendly only to fallibilism? It does seem quite odd to suggest, after all, that we can't make *any sense* of how the subject in the deceptive case has some kind of evidence for his external world claim. And it is not clear why the somewhat undeniable thought that he does have *some* kind of evidence is theoretically loaded enough to require a prior acceptance of complex philosophical views like that of fallibilism and a conjunctive theory of veridical experience. Importantly, if there is indeed a theoretically neutral way of construing the deceived subject to have some kind of evidence for his claim, then it seems it would be an equally theoretically neutral matter to go on to compare that evidence, so construed, with the evidence that the subject in the veridical case possesses. If that comparison can be made, then we can indeed, in a non-

theoretically-loaded way, ask the kind of question which forces us to make a choice between internalism and externalism.

Fighting this battle will require that I scrutinize the *plausible* ways that we can think of the subject in the hallucinatory case (and, for that matter, the veridical case) as having "some kind" of evidence for his external world claim. In order to neutralize this objection, I must show that we are *not* required to think that because these plausible ways exist it follows that the subject in the veridical case "can't tell" he isn't being systematically deceived.

We should now ask, then: *Does the subject in the hallucinatory case have some kind of evidence, of whatever quality, for his external world belief?* Here, however, we should remember that we are assuming that the New Evil Demon thought experiment is being deployed as support for internalism. Given this, we should now acknowledge that there *is* a rather straightforward internalist-style way of understanding the sense in which the two subjects have the *same* kind of evidence for their external world belief; a way that, more particularly, is the one that seems to be the most worthwhile to focus on given my own aims in this section. Namely, it is an internalist understanding that makes use of the pessimistic fallibilist's conclusions on the matter. On this view, both subjects have *poor* evidence, simply because both fail to be able to perceptually verify that the external world belief is true, which makes any belief held by either subject that the claim is (likely) true presumptuous at best. Suppose the case is such that each of the subjects, under the appropriate circumstances, forms the belief that there is a tree in front of them. It is undeniable that neither the subject in the deceptive case *nor the one in the veridical case* could, e.g., tell *just by looking* that the claim "There is a cleverly made fake tree in

front of me" is false. This, and similar facts about what they both fail to be able to perceptually verify, points directly to the way in which their evidence is of the same poor type.

However, we know by now that admitting this by itself does nothing to confirm fallibilism. Granting that a subject in a veridical case and a subject in a systematically deceptive case are of a similar evidential predicament in that both are unable to perceptually verify that the belief in question is true *just by looking*, is compatible with saying that any human subject's *best* epistemic case is one where they can nonetheless rule out all relevant cases of deception. An infallibilist would just need to insist, plausibly enough to be sure, that merely visually-based beliefs about the external world do not represent cases of a maximally favorable epistemic case for any human perceiver. Cases of forgery, holograms, the unnoticed influence of drugs, etc., (however objectively unlikely they may be) are all *rather convincing* but of course still misleading instances of visually-based beliefs. But admitting this is compatible with any perceiver's best epistemic case being one where he/she can perceptually verify the truth of certain external world beliefs, because it is compatible, for one, with claiming that only the subject in the veridical case is, properly understood, having a *perceptual* (external world regulating) experience. Thus, admitting that intuitively any conceivable perceiver is like a systematically deceived subject in that both cannot verify *just by looking* that certain external world claims are true, is quite compatible with denying that the subject in a veridical case *can't tell* that he isn't in a *systematically* deceptive case. To repeat, if he is a normal perceiver, then on an infallibilist picture of things he can tell precisely that

merely given the fact that he knows himself to be undergoing a "plain," every day waking experience.

As far as I can tell, the above commonality exhausts all that is intuitive in claiming that the two subjects have *the same kind of evidence* (according to the relevant internalist understanding). We want to say that they have the same evidence because we realize that they are indeed in an epistemic case that is of an identical type. Both cannot tell just by looking that all sorts of statements that would have it that they are not looking at the real thing are false. But that is the extent, I think, of the "intuitively obvious" lesson we can draw from the neutral formulation of the thought experiment. And, of course, an *infallibilist*, whose view pertains to the *best* epistemic case, can admit that that is the lesson without qualm. The fallibilist simply makes an illicit jump from what is indeed true of the epistemic case for any normal perceiver's merely visually-based beliefs, to a view about what is therefore their *best* epistemic case.

It may be thought, however, that the fact that it was a merely visually-based belief is incidental. A similar, neutrally-formulated thought experiment can be devised where two subjects each go through a very prolonged, arduous, and meticulous process of perceptually-based verification, but where they are similar internally-speaking yet different (in the relevant sense) externally-speaking. Suppose that the subject in the veridical case actually *does* rule out *all* cases of deception for a certain external world belief. On an infallibilist understanding of *what it takes to meet the standards for epistemic justification*, he would have an epistemically justified belief. Indeed, he would know with complete evidential certainty that it was true. But, of course, the subject in the deceptive case would, by comparison, only *think* he had gone through the relevant

process of verification and hence justification. He would be unjustified, according to an infallibilist understanding of what it takes to meet the standards of epistemic justification, because he would have not verified anything at all. Can trouble be made for the infallibilist because of the intelligibility of this thought experiment?

I don't think so. What it shows, at most, is that whether or not one is justified will depend on *more than just what happens to one from the skin inward*. But the infallibilist should claim that perceptual experience itself is *constituted* by more than just what happens inside one's own skin. He should claim that it is constituted, in part, by the proper regulative role played by the "common sense" external world. This is compatible with affirming that whether one is in a strong epistemic position with respect to external world claims is a matter determined by what one can perceptually verify, as long as one holds that to be any conceivable *perceiver* is to be the kind of thing that continuously dances with spatio-temporal objects of the appropriate size, shape, speed, hardness, etc. So, to belabor the point, drawing the above lesson from the just formulated thought experiment is compatible with saying that any perceiver *can* tell, just by knowing they are undergoing a perceptual experience, that they aren't a victim of systematic deception.

Of course, it is important to note that none of what I have said so far should be construed as a positive argument for the view that all "plain" perceptual experiences are veridical experiences. Those arguments have already been given. It is, rather, an attempt to show that there is an *infallibilist-friendly* way of explaining the intuitive sense in which the evidence of a veridical perceiver is comparable to that of a systematically deceived subject. Far from being merely dismissive of the thought experiment and its attendant conceptual apparatus, I have tried to show, first of all, that it is perfectly *legitimate* to

question whether the usual way it is formulated is coherent--the way that claims that the two subjects are undergoing identical phenomenological experiences--and that furthermore one can accommodate the core intuitive force of a more neutral formulation in strictly infallibilist-friendly terms.

I will now move to examining the case of forgotten evidence, which is a thought experiment designed to confirm externalism to the detriment of internalism.⁷³ We are asked to imagine someone who has a reliable memory, and who correctly recalls something, but who has also forgotten the ultimate evidential source of her knowledge. Suppose, to keep things simple, that she also fails to have any view whatsoever about the extent of the reliability of her own memory. A certain kind of externalist would claim that, even so, since the subject has a reliable memory she has a justified belief. It matters not to this sort of externalist that she fails to recall the basis for her belief. Given the set-up of the thought experiment, it is obvious that she would not be strongly positioned to discursively refute the claim that she is completely wrong. She would *have* to come out sounding dogmatic, it seems, if she was pressed on this matter. But so much the worse, this sort of externalist would claim, for the thought that one's ability to discursively defend oneself is something relevant to the question of whether one is epistemically justified.

Again, the problem for my position is that there seems to be a perfectly legitimate, non-theoretically-loaded question asked here, that concerns whether or not such a subject is justified. How one answers it will, once again, seem to commit one to looking at the matter in either an internalist or an externalist fashion.

It is compatible with infallibilism regarding what our epistemic position can be with respect to external world claims, however, to think that the above subject is in a *highly favorable* epistemic case *of some sort* for the case imagined. The infallibilist of this sort can admit that the subject has good quality justification *of some sort* for her belief; and just point out that this sort of case is not, they think, the best case. However, this then presents the following issue. To give this response would seem to require the infallibilist to take a somewhat guarded externalist line, and so deny internalism. If one thinks the subject in this last case is in *any* kind of favorable epistemic case, then it seems it could only be in an externalist sense--at least according to one common way of understanding the internalist/externalist controversy. An internalist, so understood, would deny that the beliefs in cases of forgotten evidence had *any* favorable epistemic status whatsoever. Dealing with this problem requires addressing how one *can* be in a highly favorable epistemic case for instances of forgotten evidence that is also a sense *neutral* with respect to the internalist/externalist controversy.

Thankfully, the relevant terminology is somewhat ready at hand. Trivially, the thought experiment has it that *we* admit/deny that the subject is justified. The subject is thus here *justified/unjustified according to how we, the one's examining the thought experiment, understand matters*. She has a belief that *we* can understand to be both true as well as reliably-based; *we* understand her to have produced the belief by memory and *we* are assuming her memory is reliable. Notice that, in what is a stark contrast, given the design of the case it is *also* true that the imagined subject is not like us in this respect. She, unlike us, fails to understand that the belief is based on a reliable process, and she certainly can't verify *that* it is a true belief. Ultimately, it is the simple fact that it is quite

consistent for *both* of these evaluations to hold simultaneously, for the case in question, which is what provides the infallibilist with an avenue for addressing this thought experiment. Under this kind of description, the intuitive result delivered by the thought experiment is, in part, the result of the somewhat obvious fact that *attributions of a belief's probabilistic relationship to the truth are perspective-relative*. It is a logical truth that there can be no attributions without an attributor. And part of what is confirmed by the thought experiment is that the same belief can, for example, be *known* to be highly probable by one party, but not known to be highly probable by another. The subject in the thought experiment clearly can be said to have a highly probable belief *according to us*, the one's puzzling over the thought experiment. But this is compatible with acknowledging that the imagined subject also does not have a highly probable belief *according to her own perspective*. For my purposes, what is important to note is how this makes the *choice* of whether one looks at this matter internalistically or externalistically, based merely upon what the thought experiment tells us, completely unnecessary. We could, for all the thought experiment tells us, adopt both ways of speaking simultaneously.

6.3 Infallibilism Defended

The goal in this concluding section is to better explain the connection between a non-cognitive theory of perception and infallibilism. It will be helpful to begin by once again asking the question with which we began: *Why be a fallibilist?* Since we've already seen that we do not *have* to accept fallibilism as a result of our inability to answer skeptically-focused questions, any answer provided here must proceed with considerations independent of its reasonableness in that regard.

Conceptually speaking fallibilism does have this much in its favor: it strikes an intuitive middle ground between opposing ways of understanding what meeting the standards for epistemic justification has got to do with being correct. On one extreme there is the rather implausible view that whether or not one is justified/unjustified is *exclusively* determined by whether the belief in question is true/false. This view is obviously problematic because the epistemic worth of a belief depends, at least in part, on certain characteristics of the process used to arrive at and/or maintain it. Lucky guesses, for example, are paradigmatic instances of unjustified beliefs. Clearly enough, then, whether or not a belief is true is not *solely* what determines its justificatory status. At the other extreme is *infallibilism with respect to meeting the standards for epistemic justification*. Infallibilists of *this sort* claim that the only epistemically justified beliefs are those that are guaranteed to be true by the process of acquisition and/or preservation upon which they are based. On this view, the question of epistemological realism versus skepticism with respect to some domain of facts is simple. We are able to *consistently* say that there is some domain of facts that we have come to know of *only* when talking about those beliefs we have concerning that domain that are evidentially certain. If a belief's truth isn't evidentially guaranteed by the process upon which it is based, on this view, then there is no circumstance under which it can be considered knowledge.

Obviously enough, infallibilism of this sort honors well the requirement that certain characteristics of the process used to form/maintain a belief are important for determining its justificatory status. But most philosophers find it to be too strong. Optimists in particular find it problematic given its failure to allow for epistemically justified external world beliefs that are nonetheless false. Suppose, for example, that an

experienced (and sober) ornithologist is staring into her garden on a clear day, and that she takes note of what is in fact a real goldfinch. One might think that *just by looking*, in this (as it happens, non-deceptive) case, the ornithologist has become epistemically justified in thinking that there is a goldfinch in her garden. But if we accept this thought, we must also accept fallibilism. Presumably even an experienced ornithologist cannot tell *just by looking* that the object in her garden isn't, for example, a highly sophisticated fake goldfinch of some variety. Thus, whether or not the ornithologist *knows* that there is a goldfinch in her garden will depend on something independent from the fact (according to the optimist) that she has an epistemically justified belief that there is a goldfinch there. She would have an epistemically justified belief even if, contrary to the case in question, there was in fact just a cleverly-made fake goldfinch in her garden. But since an infallibilist with respect to epistemic justification would have to deny this, intuitively (says the optimist) this view must be incorrect.

The infallibilist of the above sort can respond, however, by insisting that the intuitively evident point in the above example is *not* so much that an infallibilist view is too strong, but rather that there can be cases where a subject is both in a *strong epistemic position* (at least in an optimist sense) as well as misled or deceived in some fashion. The infallibilist regarding the standards for epistemic justification can agree with this completely, however, and just go on to insist that such cases nonetheless never constitute *knowledge*; even when the belief also happens to be true.

In any case, even if this is deemed to be an inadequate reply (as it very well may be) it is absolutely crucial to separate infallibilism *as a description of our actual epistemic position* with respect to certain external world claims from infallibilism as a

standard for epistemic justification for such claims. Only a defense of the former version of the view is required to address the problem of philosophical skepticism in favor of epistemological realism, and of course its plausibility is unaffected by the above considerations raised by the fallibilist regarding epistemic justification.

So, with this clarification in hand, we can now ask: what is the connection between a non-cognitive theory of perception and infallibilism so understood? The crux of the matter, it seems to me, will involve how a non-cognitive theory's truth sheds light on what it means to *perceptually distinguish* an external world belief's truth from deceptive cases of its falsehood. Prima facie, the non-cognitive theorist presents a perfectly coherent picture of perceptual experience that allows us to claim that *sometimes--perhaps, e.g., only after an arduous and meticulous process of verification--* we *are* able to establish that matters stand so that a perceptually-prompted/maintained belief is true rather than deceptively false in any way. Dewey offers something close to the relevant idea in his book *The Quest for Certainty*. He summarizes his own view on this matter, a position he calls *experimental empiricism*, as follows:

In experimental knowing, the antecedent is always the subject-matter of some experience which has its origin in natural causes, but which, not having been controlled in its occurrence, is uncertain and problematic. *Original objects of experience are produced by the natural interactions of organism and environment, and in themselves are neither sensible, conceptual nor a mixture of the two. They are precisely the qualitative material of all our ordinary untested experiences.* (pp. 172-3, my emphasis).

Since this kind of view understands perception *at the outset* as a natural interaction between organism and environment, I would claim (unfortunately, for reasons I will not get into here, *pace* Dewey) it is the sort of view that has it that a perceptually-based process can be employed to effectively rule out all problematic epistemic

possibilities, the obtainment of which depends on mind-independent reality being arranged in a manner contrary to the truth conditions for a certain belief about one's immediate surroundings. It is friendly to the idea that after we have, for example, prodded the relevant dry-goods enough, we are able to establish that the belief is true *simpliciter*, and not deceptively false in any imaginable way.

Explaining how such a process of conclusive verification might work allows me, as it turns out, to integrate what was shown above to be the plausible aspects of pessimistic fallibilism and the various forms of non-discursive optimism examined in chapter 3. What we can take from pessimistic fallibilism is the view that we can achieve a certain kind of strong epistemic position by effectively answering certain kinds of skeptically-minded questions. This is precisely what knowing we have successfully performed an act of conclusive verification would accomplish. What Wittgenstein and Quine have taught us, albeit in starkly different ways, is that epistemological beliefs require non-epistemological beliefs. In our case, the crucial non-epistemological belief is the acceptance of a non-cognitive theory of perception. It is when this claim operates in the background as a "hinge" proposition--as what helps to determine the character of an employed logic of inquiry--that, I hope to show, we are allowed to comfortably say that a certain perceptually-based process of conclusive verification is, in fact, successful. Process reliabilism taught us that both a claim's high objective probability as well as the role that it plays in our cognitive economy are crucial for determining our positive epistemic position with respect to it, at least in the context of skeptically-minded questioning regarding our basis for asserting that claim. Being known to be a conclusively verified claim honors both of these constraints quite straightforwardly.

Finally, contextualism taught us that we must *first* determine our epistemic position before we can properly assess whether or not that position is in fact strong enough to *allow us* to effectively answer certain kinds of skeptically-focused questions. The net effect of the discussion to follow, or so it is hoped, is to persuasively suggest that the truth of a non-cognitive theory entails that our epistemic position is in fact strong enough to at least allow for the conclusive verification of certain perceptually-based claims about external reality.

We can begin to draw all of these considerations together under one coherent picture when we realize that the known truth of a non-cognitive theory allows us to say that our epistemic position is, in fact, strong enough to allow us to rule out *the obtainment of skeptical scenarios*. But, of course, to know that one is not dreaming, a brain in a vat, etc. is not yet, at least in every attainable case, to know this or that about one's external surroundings. The crucial explanatory gap that needs to be filled, then, is why our known epistemic position regarding the skeptical scenarios allows us to know that we can conclusively verify this or that about the external world. Thankfully, as far as attacking this otherwise potentially overwhelming issue is concerned, we are helpfully constrained in the following sense. Since it is the known truth of a non-cognitive theory of perception which allows us to determine our epistemic position regarding the skeptical scenarios, any considerations offered to fill this gap must fall out of what a non-cognitive theory demands.

According to a non-cognitive theory, then, perceptual experience is characterized as a certain kind of *continuous, perceptual-experience-regulated interaction* between a subject and their surrounding environment. Baldly stated, the nature of that interaction is

best understood as that of a kind of behaviorally-mediated feed-back loop involving two different types of events: (i) a typical human subject's perceptual experiences; and (ii) the arrangement of middle-sized physical objects in their immediate surroundings. Crucially, according to a non-cognitive picture of things, it is completely misguided to treat the relation here as analogous to that of chicken and egg, inviting the devilish question of which comes first. Instead, the relation is best seen as exemplifying, in what is simply a more generically-applicable sense, the same features as the following relationship: that of two people dancing with one another where neither is, in fact, leading. According to a non-cognitive theory, we are continuously and quite literally *dancing* with middle-sized external world objects in our immediate surroundings, and, although we often like to *think* otherwise, it is ultimately most accurate to maintain that neither "party" is leading (an immediate implication of this last point might even turn out to be the non-cognitive theorist's way of describing the essence of the problem of induction).

Of course, the above somewhat playful remarks raise more questions than they answer. Indeed, there is much more that needs to be examined and clarified here regarding the non-cognitive theorist's understanding of the character of our relation to external world objects, and what one might call the *sensorimotor* account of perception that falls out of it. Also, it goes without saying that this general sort of view of perceptual experience is nothing new.⁷⁴ But at this point I would like to turn to the following crucial *epistemological* issue: to know that a non-cognitive theory is true, I would contend, is to know that certain perceptually-guided behavioral interactions with one's surroundings are related to certain beliefs *as a successful conclusive verification procedure for a certain claim is related to a claim so verified*. How to properly answer

the question (assuming, that is, that there is even a single, determinate answer here) of what *general* criteria ought to be employed for determining which verification procedures and claims are so related is an extremely complicated matter that I cannot adequately address here. But I do think I can here present plausible enough considerations for why the known truth of a non-cognitive theory allows for one to know that *some* such verification procedures/external world claims will, in certain cases, have to be so related.

As a way of properly entering the exploration of this matter, we should take note of the following distinction. If we are after an explanation of how a non-cognitive view of perception leads to infallibly justified beliefs about the external world, then it is best to do so, I think, by keeping the following quite distinct kinds of examination separate. On the one hand, we could examine what I would call the *contingently historical, subject-dependent* facts that allow for this kind of belief; and, on the other, we could examine the *contingently present, subject-independent* facts that do so.

I have already, in fact, touched upon what I think are some of the important examples of the former set of facts, and I will have more to offer in that regard in just a moment. But the best way to enter that discussion, I believe, is to first say some things about the other set of facts; the contingently present, *subject-independent* ones. To that end, we can say that to know that a non-cognitive theory is true just is *to know that* we are in continuous perceptually-regulated contact *with a mind-independent reality that is stable enough to allow for the truth-conditions of certain utterances to be met for considerable enough periods of time, and in a manner that is not affected in any way whatsoever by the mere opinions, wishes, biases of human subjects*. It is to know that I am in contact with a mind-independent, spatio-temporal reality the arrangement of which

will not be altered by a mere change in my manner of thinking; that it is a mind-independent reality that is more stable than that, and that for this reason the truth-conditions of certain utterances will be met in a way that will stay equally rock-solid (or, as the case may be, not). Knowing that a non-cognitive theory is true allows me, for example, to know that I'm in contact with a mind-independent reality that is stable enough to have allowed for the truth conditions of the utterance "There is a coffee cup on my table" to be met, for the most part, for the duration of this morning. But, importantly, the very statement of this realization raises the following issue; and it is in this way that the discussion of the admittedly obvious subject-independent facts that allow for infallibly justified beliefs about the external world leads to a further discussion of subject-dependent ones. Namely, in the above statement I mention *an utterance*. And, of course, here I am talking about a *cognitively significant* utterance, since only this sort of utterance will have truth conditions.

It seems clear enough, then, that to fill the explanatory gap of explaining how a non-cognitive theory leads to infallibly justified external world beliefs, we must, not surprisingly, confront the thorny matter of meaning. We must examine, that is, what a particular subject (rather than some other subject with a different personal history, physiology, etc.) *means* by the declaration of a sentence that they are willing to assert outright in a case where, the non-cognitive theorist hopes to show, they know they are in fact infallibly justified. The precise reason this matter is of central relevance for filling the relevant explanatory gap, is due to the obvious fact that an asserted sentence's meaning can be context/linguistic community relative, in such a way that, depending upon how matters stand, it may or may not be able to be shown to be infallibly supported

in a particular case. To employ a simple example, suppose that in a certain context you (or your linguistic community) mean something different by the sentence "The cat is on the mat" than do I (or my linguistic community). It might follow, of course, that we consequently disagree over whether a certain procedure serves to conclusively verify the truth of that claim. What needs explaining, then, is how the fact that disparities such as this are bound to arise is consistent with the claim that even in such cases the establishment of flawless justification can nonetheless at least sometimes occur. Only a careful study of what goes on in such cases (or at least what can go on) can reveal if infallibilism still comes out looking promising along non-cognitive theorist lines. Here the contextualist will, I think, have much to offer that is helpful. But a proper examination of this tricky matter must wait for another time.

That said, although I think they are far from conclusive, let me offer the following remarks on behalf of the infallibilist. Namely, I would like to make use of a key contingently historical, subject-related fact that has been established as a result of knowing that a non-cognitive theory is true. Anyone who is able to establish the truth of a non-cognitive theory in the manner laid out in the previous chapter, can say that they have established that they themselves are a human subject of the following sort: one who knows that what they take to be normal, everyday waking experience is in fact waking experience in which they are involved in behaviorally-regulated commerce with a mind-independent spatio-temporal reality. In order to see why it appears plausible to think that the explanatory gap just mentioned can be filled, then, it will be helpful to emphasize certain implications of the sort of *contingency* that is an inextricable part of being one who can establish this about oneself. Ultimately I can only speak for myself when I say

that I'm able to follow the logic of the previous chapter in a way that makes a non-cognitive theory come out looking rather convincing. There is no doubt that the reasoning could be helpfully developed in various ways, but even in the above form the arguments seem, at least on the whole, rather sound to me. But of course, on the one hand, there will be readers who also have carefully considered the above arguments and yet do not see why a non-cognitive theory is so obvious (perhaps as a result of feeling the force of an even higher-order form of skepticism that I have yet to get a handle on); and, on the other hand, it is equally obvious that I myself may lose confidence in that theory at some later point. But it is crucial to note that this does not affect the fact that there will be those who are able at a certain time--as a result, to be sure, of contingently historical, subject-related facts about themselves--to follow the above reasoning and thereby become *soundly* convinced of a non-cognitive theory's truth. Consider an analogy. A person with average vision is able to tell from a short distance that a certain written letter is, for example, a "T" rather than an "I." Obviously, however, that person may lose that ability with time, and at any given time there will be those who lack that same ability. But equally obvious is that this does not affect the fact that at a certain time the person with average vision will be able to so *reliably* and *sensitively* discern how matters *actually* stand. I think it is correct to look at things in at least a roughly similar way for the one who is able to understand that a non-cognitive theory is true as a result of the line of reasoning presented in the previous chapters. The key difference between these two different kinds of subject is just as follows. The non-philosophical person with normal vision has the matter of what constitutes a verification procedure *fixed naturally* for them, so that they just proceed by way of habit with using their sense experiences to verify reliably whether

or not a certain letter is a "T" versus an "I." While the philosopher who has come to know that a non-cognitive theory is true in the above-described sense is one who has the matter of what constitutes a verification procedure for a non-cognitive theory be something that, they can openly admit, is *fixed somewhat on-the-fly* out of their desire to avoid the paradox that would arise from having to accept fallibilism. It is the desire to avoid this paradox, in other words, that constitutes the normative force that drives their reasoning. And, of course, the reason the outcome of this fixed "on-the-fly" methodology has relevance to contemporary epistemology--the reason that the artificial-looking methodology is nonetheless something the outcome of which philosophers should find interesting--is that the fix was seen as *necessary* in order to make what philosophers like Stroud claim Descartes' methodology in the *Meditations* is look remotely philosophically respectable.

With this I can begin to gesture towards why I am optimistic that the above explanatory gap can be filled in a way that makes infallibilism remain plausible. Because of the fact that those who possess the relevant ability know, however contingently, that they are interacting with a stable enough, mind-independent spatial-temporal world, it *seems* that the only thing that would prevent them from going on to claim that, in certain cases, they know that a certain procedure was successful in establishing the truth of a certain claim about the external world, was the protestations to the contrary offered by another subject. But, given their knowledge, in this case, of the important difference between the kind of relationship that perceptual experience affords them with external reality and the kind of relationship that they might bear to a different subject (with a different personal history and, perhaps, physiology) it should be equally clear to a subject

that knows that these two kinds of relationship hold that there might be certain conversational contexts where properly resolving a disagreement with someone else regarding whether a claim is conclusively verified by a certain procedure will be impossible. But, I would contend (predict), for such a subject, realizing this is compatible with being assured that he/she possesses the relevant knowledge concerning mind-independent *and* conversation-independent reality. Ultimately, this is because knowing that a non-cognitive theory is true allows us to recognize that we have a proper epistemological basis for the following claim: one of the things we do is speak a language--and, therefore, sometimes, ask/answer epistemological and meta-epistemological questions--another is exist in a spatio-temporal universe that is not of our own making. In sum, then, the reason that I am at least optimistic that the above explanatory gap will be filled properly for the infallibilist is that the verbal nature of the above type of disagreement seems to assure that either (i) proper steps can be taken to make it so that, at least in certain cases, soundly based agreement that a claim is conclusively verified can eventually be reached between the subjects; or else that (ii) the impossibility, in certain cases, of reaching that sort of agreement in practice presents no epistemological issue (given the just-mentioned truism that falls out of knowing that a non-cognitive theory of perception is correct).

Of course, the above remarks in no way establish (i) and (ii). But all I am trying to do here is gesture towards how I think the matter can be worked out. A full examination of the issue would have to deal with the problems that, to be sure, appear to arise on all sides. Here I will mention just two. First, a legitimate worry surrounds whether a non-cognitive theorist is forced to accept a problematic form of relativism. How else, it might

be asked, can she live with the acknowledged possibility of the above sort of intractable disagreement? Secondly, there is the much more easily anticipated worry that a non-cognitive theorist who arrives at their view in the manner laid out in the previous chapter must sound dogmatic according to following kind of disputant: a fittingly characterized *third-order skeptic*. How can I be so assured that the above type of examination conclusively guarantees the truth of a non-cognitive theory in a way that ably refutes a second-order skeptic? Is there not *some* way to cast legitimate-sounding doubt on this claim? If nothing else, my confidence here seems, once again, to fail to properly acknowledge that definitive refutation in philosophy is a rare feat indeed.

Here I cannot give these questions the attention they admittedly deserve. Nevertheless, I can say the following regarding the second worry. It is crucial to note that just as the committed fallibilist (conjunctive theorist of veridical experience) is not a second-order skeptic, so too is it correct that the second-order skeptic is not a third-order skeptic. To repeat, the arguments of the previous chapter were designed to convince someone who did not treat it as immediately obvious that they do not know they are not dreaming, and still wondered whether they should accept second-order skepticism. They were designed, more accurately, to refute the second-order skeptic, where what position the second-order skeptic maintained was a matter that was rather narrowly conceived. Most notably of all, it was understood that the second-order skeptic was one who agreed that a properly philosophical examination of our knowledge could reach conclusive results. The second-order skeptic believes that that examination conclusively reveals the following: the only way to maintain the claim that one knows something about external reality is to beg the question against the claim that one does not, and vice versa. In

contrast, I have maintained that that examination conclusively reveals that a non-cognitive theory is correct. But each party here agrees that the examination can conclusively reveal *something*. It is here, of course, where the third-order skeptic will try to make headway. But the epistemological situation at this point, absent a convincing argument from the third-order skeptic, is peculiar; peculiar, to be sure, in a way that is to be expected when one tries to *assert* a substantive philosophical view. On the one hand, it is certainly incorrect to say that one who has refuted the second-order skeptic in the way described in the previous chapter thereby knows that a third-order skeptic has to be wrong. But until what *it is* to be a third-order skeptic is precisely defined; until that sort of skeptic makes a persuasive case for their view, it is equally *incorrect* to say that the mere gestured-towards possibility of third-order skepticism shows why one who stands where the non-cognitive theorist now stands is uncertain. Consider the following analogy. The fallibilist would never be able to *persuasively* make her case were she not able to present the conceivability of the skeptical scenarios as the reason why we do not know with complete evidential certainty anything about external reality. Until such reasons are provided, we obviously have no basis for being *convinced* that we do not know this or that about external reality. If all the fallibilist did was stamp her feet and insist that we do not know what we say we know about external reality, then the best response would obviously be to simply walk away.

For this same sort of reason, the bare assertion of third-order skepticism is, at least for right now, of little moment. This is not to say that a forceful version of it will never be able to be motivated. And if at some point someone clever enough does persuasively motivate third-order skepticism, then the same general mismatch problem will arise; what

we take our epistemic position to be with respect to external world claims would be demonstrably inconsistent with what our acknowledgement that we don't know that we know that we know this or that about external reality entails it to be. If the third-order skeptic presents considerations so forceful that they *show* that we can't refute the claim that the examination of the previous chapter failed to be conclusive in supporting a non-cognitive theory, then we will be unable to henceforth maintain that we have, in fact, conclusively verified that theory. And if our confidence in the conclusive verification of a non-cognitive theory disappears, then likewise gone is our basis for maintaining that we know anything about external reality. Nevertheless, for phenomenological reasons *that will not go away as long as we are built the way that we are now built*, we will *still* also be forced to think that our epistemic position with respect to external world claims is as strong as ever. In this (at the moment quite fanciful) instance, the problem of philosophical skepticism would present itself anew. But if it did, so be it. Such is our lot in life. *If* third-order skepticism can be effectively motivated, all it would show is that the search for *philosophical* satisfaction is, not surprisingly, like any other; it is (unfortunately?) naïve to think that one's work will ever be done.

NOTES

¹ See Stroud (2009) for discussion.

² That an intimate connection exists between one's metaphysics of perception and one's epistemology is perhaps made most evident by the very existence of the rich subfield in philosophy known as *the epistemology of perception*. Philosophers interested in the epistemology of perception ask, among other things: what follows regarding what do we know and/or justifiably believe given what is evidently true regarding perceptual experience? The following is a list of helpful readings that relate directly to central issues in the epistemology of perception: Alston (1993); Armstrong (1961, 1968); Austin (1962); Ayer (1946-7); Block (1996, 1997, 2004); BonJour (2004); BonJour and Sosa (2003); Brewer (2000); Burge (1982, 1991); Chisholm (1957); Crane (1992, 1998, 2001); Crane and Farkas (2004); Dancy (1988, 1995); Davies (1992); Dretske (1969, 1981, 1995); Ducasse (1942); Foster (2000); Fumerton (1985); Goldman (1986); Grice (1961); Hinton (1973); Huemer (2001); Hume (1748); Jackson (1977); Johnston (2004); Loar (2002); Lycan (1996, 1999); Martin (1992a, 1992b, 1998, 2002); McCulloch (1995); McDermid (1994a, 1994b); McGinn (1989); Merleau-Ponty (1945); Millikan (1987); Noe and Thompson (2002); Peacocke (1983, 1992); Perkins (1983); Pitcher (1970); Price (1932); Pryor (2000); Russell (1912); Shoemaker (1990); Smith (2002); Snowdon (1990, 1992); Spener (2003); Strawson (1979); Sturgeon (1998); Swartz (1965); Thau (2002); Tye (1984, 1992, 1995, 2000); Valberg (1992);

³ See Alston (1986, 79) for discussion.

⁴ The more familiar sense in question relates to how beliefs about what one knows can inform beliefs about the nature of perceptual experience. For example, if one believes that one can only non-inferentially know about the status of one's experiential states, then one would have to deny a theory of perception such as direct realism, etc.

⁵ See, for example, Stroud (1984, pp. 132, 158, 179-180, 194-195, 214, 256), and (2000, pp. 8, 29-31, 47-50, 56-58, 97, 106, 120, 133, 154). See also Fumerton (1995).

⁶ See Byrne (2004); Clarke (1972); Fumerton (1995); Greco (2000); C. McGinn (1989, 113-4); M. McGinn (1989); Nagel (1986, 67-74); Stroud (1984); Williams (1996); and Wright (1985).

⁷ A popular way that a conjunctive theory is realized these days is in the claim that perceptual experience is constituted, in some sense, by *representational contents*. See, e.g., Armstrong (1968); Block (1990); Burge (1986); Byrne (2009); Chalmers (2006); McGinn (1989); Peacocke (1983, 1992); Pautz (2010); Searle (1983); Shoemaker (1990, 1994); Siegel (2010b, 2010c). For general discussion see Siegel (2010a).

⁸ At least regarding its focus on how we are *already situated* in a spatio-temporal world, this position seems somewhat close to views about perceptual experience held by certain 20th century European phenomenologists; notably Heidegger, and Merleau-Ponty. Dewey also expresses views that I think are friendly to a non-cognitive theory; see, for example, Dewey (1917, 64, 71); (1929, 172-173). As far as contemporary versions of the view go, see: Noë (2002a, 2002b, 2002c, 2004, and 2006) for discussion of a roughly equivalent idea. It is not clear to me, however, whether Noë would endorse a non-cognitive theory of perception in particular as I present the position in this work.

⁹ For helpful discussions of fallibilism see Feldman (1981); Cohen (1988).

¹⁰ Discussions of issues central to the internalist/externalism controversy can be found in: Alston (1986a, 1988, 1989, 1995, 2005); Armstrong (1973); Bach (1985); Bergmann (2006); BonJour (1980, 1985); Chisholm (1977, 1988); Cohen (1984); Conee and Feldman (2001, 2004a, 2004b); Feldman (2004a, 2004b, 2005a, 2005b); Foley (1993); Fumerton (1988, 1995, 2004, 2006); Gettier (1963); Ginet (1975); Goldman (1967, 1979, 1980, 1988, 1993, 1999, 2001); Kornblith (1988, 2001); Lehrer and Cohen (1983); Nozick (1981); Plantinga (1993); Sosa (1991a, 1991b); BonJour and Sosa (2003); Prichard (1950); Steup (1999, 2001); and Stroud (1989, 1994).

¹¹ Contemporary philosophers who, to varying extents, endorse the reasonability of philosophical skepticism include: Fumerton (1995); M. McGinn (1989); Nagel (1986); Stroud (1984, 2000); Unger (1975).

¹² For Descartes's original discussion of the dream argument see: Descartes (1641). Worthwhile contemporary discussions of Descartes's arguments, and his meta-epistemology in general, include: Adams (1975); Bennett (1990); Beyssade (1993); Bouwsma (1949); Broughton (2002); Carriero (2009); Chappell (1986); Chisholm (1982); Cuning (2007); Curley (1978, 1986, 1993); Della Rocca (2005); DeRose (1992a); Doney (1955, 1987); Dunlop (1977); Frankfurt (1970); Friedman (1997); Garber (1986, 1992); Gaukroger (1989); Gewirth (1941); Hacking (1980); Hintikka (1962a, 1978); Hoffman (1996); Jolley (1990); Kenny (1968); Lennon (2008); Loeb (1992); Markie (1992); Menn (1998); Morris (1973); Nadler (2006); Nelson (1997); Newman (1994, 2004, 2006, 2007, 2009); Newman and Nelson (1999); Nolan and Nelson (2006); Popkin (1979); Rickless (2005); Russell (1945); Sarkar (2003); Sosa (1997a, 1997b); Van Cleve (1979); Vinci (19898); B. Williams (1978, 1983); M. Williams (1986, 1996); and Wilson (1978).

¹³ However, the so-called *contextualist* about knowledge semantics will try to ameliorate the force of the result. I will address this matter further in the third chapter.

¹⁴ This characterization of a problematic epistemic possibility resembles one way that one may characterize a so-called *defeater*. Talk of defeaters in the literature, (unfortunately) like many other philosophical notions, is terribly unregulated; it can mean a slew of different things depending on who you read. This is why I have avoided using the term, although I suppose I could have chosen to call problematic epistemic possibilities, for example, something like *potential propositional defeaters*. Propositional defeaters, generally speaking, are true propositions that prevent a justified (according to an *intuitive* understanding of what it means to be justified) true belief from counting as knowledge. I would have to include the term "potential," of course, because propositions known to not be known to be false may nonetheless be false. Nevertheless, a problematic epistemic possibility does resemble a propositional defeater in that if it were added to the subject's evidence base, it would make it so that some belief that p failed to be adequately supported by one's overall evidence base (here leaving open whether or not the belief is adequately supported when the evidence base does not include the propositional defeater/problematic epistemic possibility). Propositional defeaters are also sometimes called *knowledge defeaters*, *factual defeaters* and even *overriding defeaters* (phew!) because they are propositions *the mere truth of which*, to repeat, supposedly prevents an (intuitively) justified true belief from counting as knowledge. Discussion of propositional defeaters began as a result of Edmund Gettier's (1963) paper which tried to show that

justified true belief analyses of knowledge are subject to clear counter-examples. The literature surrounding this issue is simply enormous. See Shope (1983) for an introductory overview. Knowledge defeaters are contrasted with so-called *justification defeaters*. This kind of defeater, also sometimes called a *mental state defeater* or *overrider*--perhaps one can now see what I mean when I say that discussion of defeaters is highly unregulated--is, according to its most common use, supposed to prevent a *previously* justified belief from subsequently counting as justified. These are mental states (such as believing something) which make it so that one's justification for a certain belief is rendered null. Now, if we extend this notion, as some philosophers do, so that it includes mental states which prevent some beliefs from ever being justified *in the first place*, (or which showed that these beliefs were never justified at any point whatsoever) then one accurate way to describe the arguments that comprise the gist of the next three chapters is as attempting to prove that the skeptical scenarios are problematic epistemic possibilities which are also justification defeaters in this last sense. The whole point of the arguments for skepticism is that the upshot is supposed to be that one loses all confidence in the claim that one's external world beliefs are justified--clearly a kind of "defeat." However, since I think central talk of justification defeaters involves their function of *taking away* justification that was there previously, I think it is best to also avoid calling problematic epistemic possibilities defeaters of this sort. Problematic epistemic possibilities can cause us to lose confidence in our external world claims even if those claims were *never* justified in the first place. There is one more distinction drawn between kinds of defeaters that is worth noting. So-called *undermining defeaters* are those which merely serve to nullify the support for a certain belief that p, without actually providing support for not p. They are considerations which are designed to show that one's support for a certain claim is inadequate, according to some previously chosen standard for adequacy. *Rebutting defeaters* are, in contrast, considerations designed to show that the negation of a certain claim is what is, in fact, most reasonable to maintain. This distinction was introduced by Pollock in his (1970); it cuts across the distinction between knowledge and justification defeaters. Of course, to once again state the obvious, problematic epistemic possibilities are supposed to be, if anything, a kind of undermining justification defeater for the claim that one is *absolutely certain* of the truth of a certain claim. They are supposed to undermine one's confidence in the undeniable strength of one's evidence for a certain claim. But whether they *nullify* one's justification *altogether* for a certain belief is, of course, debatable. Discursive optimists think that they don't; pessimistic fallibilist's think they obviously do. Since we are, at least for the sake of argument, leaving this question open at the moment, it is best, I think, if I refrain from calling problematic epistemic possibilities undermining defeaters of any sort. What is more, since a discursive optimist is a fallibilist, it would be misleading to consider problematic epistemic possibilities undermining of anything at all, at least according to their view. According to discursive optimists, it is simply impossible to ever be *absolutely certain* of the truth of external world claims. The existence of problematic epistemic possibilities merely at best points this out, which is not really to undermine anything at all. For helpful discussions of defeaters see Alston (1986, 1988a, 1988b, 1989, 2002, 2005); Annis (1973); Barker (1976); Bergmann (1997a, 1997b, 2000, 2005); Chisholm (1989); Goldman (1976, 1986); Harman (1973); Janvid (2008); Klein (1971, 1976, 1981); Lehrer and Paxson (1969); Nozick (1981); Pollock (1970, 1974, 1984, 1986); Shope (1983); Steup (1996); Swain (1974, 1981); and Swinburne (2001).

¹⁵ To require that a belief not be true "accidentally" is related to imposing a so-called *safety condition* on knowledge. According to a safety condition, in a case where one believes truly that p, a necessary condition for knowing that p is that in all roughly

similar situations where the subject believes that p on the same basis, p is true. Recent discussion of the requirement that beliefs not be true "by accident" in order to be considered knowledge can be found in Pritchard (2005a). For further discussion see: Coffman (2007); DeRose (1995); Goldberg (2007); Greco (2003, 2007); Hiller and Neta (2007); Nozick (1981); Pritchard (2003, 2005b).

- 16
1. $P \rightarrow Q$ [premise]
 2. $\sim Q$ [premise]
 3. $\sim Q \rightarrow \sim P$ [1, contraposition]
 4. $\sim P$ [3, 2 modus ponens]
- 17
1. $P \rightarrow Q$ [premise]
 2. P [premise]
 3. $\sim Q \rightarrow \sim P$ [1, contraposition]
 4. $\sim\sim Q$ [(3, 2 (after double negation), modus tollens)]
 5. Q [4, double negation]

18 See Carroll (1895) for the original argument.

19 Stroud (1984, 12).

20 Given his commitments to foundationalism and his views about what sort of beliefs are non-inferentially justified, Fumerton think this is the *only* way to avoid skepticism. See Fumerton (1995).

21 Classic defenses of a frequentist position, which include attempts to address the single-case problem, can be found in Venn (1876); Reichenbach (1949); and von Mises (1957).

22 Clearly enough it is not *necessary* for an infinite sequence to converge on a limiting relative frequency of a certain outcome. Take, for example, the following sequence:

010100110000111100000000*11111111**...

If the relative frequency of 1s is measured at point * then it comes out to $1/3$. If it is measured at point **, however, it comes out to $1/2$. It is clear that the noticeable pattern here can simply repeat, making the relative frequency of 1s fail to converge on a single value in the specific case where the sequence extends to infinity. It is for this reason that frequentists like, for example, Hans Reichenbach include the claim that an infinite sequence converges to a certain limiting relative frequency as an *assumption*. See, for example, Reichenbach (1949, 69).

²³ Defenses of a propensity interpretation of probability can be found in: Fetzer (1982, 1983); Giere (1973); Gillies (2000); Hacking (1965); Miller (1994); and Popper (1957). Not all of them, however, explicitly endorse tying the propensity of an outcome to the truth of a counterfactual.

²⁴ Hacking (1965), Gillies (2000), and Popper (1957), for example, would *not* go so far as to *identify* the relevant probability with a certain tendency towards a result. Fetzer (1982, 1983), Geire (1973), and Miller (1994), however, do think that we should identify the probability that a certain event has of obtaining with the relevant tendency.

²⁵ The proper understanding of counterfactual conditionals is a notoriously thorny issue. The most famous (and, at least in terms of its metaphysics, extremely peculiar) treatment is given by David Lewis (1973).

²⁶ For different criticisms of Wright's argument, see Brueckner (1992); and O'Hara (1993).

²⁷ Historically speaking, the closest view to pessimistic fallibilism is perhaps classical foundationalism. Due to a line of reasoning importantly similar to what was deployed in the previous chapter, the typical classical foundationalist is convinced that fallibilism breeds what is called *radical skepticism*--which is the view that the most reasonable philosophical position to hold is that our external world beliefs are unjustified or unreasonable. But there are well-known theoretical problems with foundationalism, and one of the goals of this work is to show that there is no reason to address these particular issues if all one is after is an explanation of what kind of general philosophical outlook results from finding the above minimalist arguments convincing. For discussion of important issues pertaining to foundationalism, including the viability of popular rival positions, see: Armstrong (1973); Ayer (1956) Bonjour (1985, 2001a); Fales (1996); Fumerton (1995); Goldman (1979, 1986, 1988); Huemer (2002); Klein (1998); Lehrer (1974); Moser (1989); Price (1932); Russell (1910-11, 1948, 1984); Sellars (1963); and Williamson (2000).

²⁸ See, for example, discussions pertaining to the debate between epistemological internalists and externalists. Relevant sources are cited in n. 10 above. Skepticism over the viability of analyzing knowledge is famously voiced in Williamson (2001, Ch. 1). Lastly, the most trenchant criticisms of the viability of philosophically analyzing any interesting concept can be found in Wittgenstein (1953).

²⁹ See Alston (1986, 79) for discussion.

³⁰ The cases in question are those where one's sense experiences alone can non-inferentially justify a relevant probability judgment about contingent reality. Defense of this position can be found in Fales (1996).

³¹ For helpful discussion of the relevant regress for strong forms of epistemological internalism see Bonjour (1985, 32); Fumerton (1988, 448); Hetherington (1990) and (1991, 860); and Steup (1989). Recently, discussion of this issue has gained prominence as a result of the work of Michael Bergmann; see especially his (2006, 9, 13-14).

³² This condition relates to the so-called *sensitivity* condition that certain philosophers place on knowledge. On this view a necessary condition for knowing that p is that one's

belief in p is sensitive to p 's falsehood; where this means, roughly, that one fails to believe that p in all situations otherwise similar to those where p holds, but such that $\sim p$. Condition (iii) exemplifies sensitivity, then, by requiring that a subject fail to believe that they are having a red sensation in a situation where everything else is the same except that, for example, they are looking at a patch of blue fabric instead of a patch of red fabric. Nozick (1981) brought this condition to prominence. It also plays a central role in DeRose's brand of contextualism (I discuss contextualism the next chapter). See DeRose (1995). For general discussion of the sensitivity condition see Black and Murphy (2007); Goldman (1976); Rott (2004); Vogel (2007); and Williamson (2000, Ch. 7).

33 The distinction between occasion and standing sentences is owed to Quine (1960, 35-40).

34 To claim that knowledge is closed under known entailment is to maintain that when one knows that p and one knows that p implies that q , it follows that one knows that q . Certain philosophers deny closure, because it seems to lead straight to skepticism. This is because knowing things about the external world clearly implies knowing that we are not dreaming, brains in vats, etc. However, philosophers who deny closure maintain the fallibilist position that we do not know the latter. Denying closure is thus the only way for these philosophers to admit that we do not know we are not dreaming, etc., while preserving the claim that we know things about external reality. See, for example: Dretske (1970, 2005) and Nozick (1981) (although in Nozick's case the reason that closure is denied is *not* primarily out of a desire to avoid skepticism, but instead has to do with what he takes to be the right analysis of knowledge).

35 For discussion of the distinction between propositional and doxastic justification (although sometimes the terminology employed is different) see: Alston (1985, 190) and (2005, 90); Conee and Feldman (1985, Section 4); Feldman (2002, 46) and (2004c, section 1); Korcz (2000, 525-526); Kvanvig (2003, section B1); Pappas (1979); Pollock and Cruz (1999, 35-36); Pollock (1986, 36-37); and Swain (1979, 25).

36 The most famous foundationalist is Descartes. Contemporary defenses of a Cartesian-esque form of foundationalism can be found in Fales (1996); Fumerton (1995); Bonjour (2001a, 2001b, 2003a, 2003b); and McGrew et. al. (2007). Other defenses of various brands of foundationalism can also be found in Armstrong (1973); Chisholm (1979); Goldman (1979, 1986, 1988); Huemer (2001); Pryor (2000); Russell (1910-11, 1948, 1984); Williamson (2000).

37 The most famous arguments against foundationalism include: Bonjour (1985, Chs. 2-4) Chisholm (1942) (in this work Chisholm deals with a problem for foundationalism--known as *the problem of the speckled hen*--the raising of which he credits to Gilbert Ryle); Davidson (1989); Rorty (1979); and Sellars (1963). More recent criticisms have come from Williams (1996). For a helpful overview of the character of the debate between foundationalists and coherentists see Sosa (1980).

38 Famous defenses of coherentism include: Bonjour (1985); Davidson (1989); Lehrer (1974); and Sellars (1963).

39 See, for example, Bonjour (2001b, 79-80); Fumerton (2001, 70); and Stroud (2000, 122-123).

⁴⁰All citations will refer to Wittgenstein (1969). For discussion of Wittgenstein's *On Certainty* see: Moyal-Sharrock (2000, 2003, 2004, 2007); Moyal-Sharrock and Brenner (2005); Pritchard (2001, 2005c); Rudd (2005); Stroll (1994) Vintiadis (2006).

⁴¹ Worthwhile discussions of the contextualist response to skepticism include: Cohen (1988, 1998, 1999); DeRose (1992b, 1995, 2002); Hambourger (1987); Hawthorne (2004); Lewis (1996); Neta (2003); Rysiew (2001); Schiffer (1996); Stanley (2005); and Unger (1984, 1986).

⁴² It should be noted that in this excerpt DeRose does not explicitly compare the claims that I am not dreaming and that I have hands. Instead of the former claim, he chooses the claim that I am not a brain in a vat. But, since each of the relevant claims here involve the negation of a skeptical scenario involving systematic deception, this switch is quite harmless. For the sake of consistency with what I have written in the previous chapters, I have chosen to frame things in terms of the claim that I'm not dreaming.

⁴³ The term "Moorean" is used here because the claim that he knows that he has hands is one among others that G. E. Moore famously and unapologetically employed to explain why he knew the skeptical scenarios failed to obtain; see especially Moore (1939).

⁴⁴ Works of Quine's that discuss central aspects of his naturalized epistemology include: Quine (1960; 1969, Chs. 3 and 5; 1974, pp. 1-4 and 137-141; 1975; 1976, Chs. 22-24; 1980, Ch. 2; 1981a, Chs. 1, 2 and 4; 1981b; and 1982, pp. 1-5).

⁴⁵ See Stroud (1984, pp. 132, 158, 179-180, 194-195, 214, 256), and (2000, pp. 8, 29-31, 47-50, 56-58, 97, 106, 120, 133, 154).

⁴⁶ Where by this it is meant an assessment that brackets considerations having to do with merely what it is appropriate to *say* we know. Stroud thinks that there can be plenty of claims that are appropriate to say we know, even though we do not in fact know those same claims to hold. For an extended discussion of the requirement for detachment, and how it affects the relevance that ordinary language philosophy has for addressing skepticism, see Stroud (1984, ch. 2). For convincing criticisms of Stroud's view on this matter see Pappas (1991) and Leite (2004). For the classic discussion of how ordinary language philosophy can helpfully address the problem of skepticism see Austin (1961). A useful general discussion can also be found in Williams (1996, pp. 15, 65-66, 172-175, 181-185, 193-194, 357).

⁴⁷ Stroud never, to my knowledge, explicitly describes what he is up to in precisely this *condensed* manner. The characterization offered above is owed in particular to Michael Williams' discussion of Stroud's views in (Williams 1996).

⁴⁸ Stroud (2000, 122-123).

⁴⁹ Stroud argues that it is part of our intuitive concept of knowledge that in order to know that p, we must know to be false all scenarios known to be incompatible with knowing that p (Stroud 1984, Ch. 2). This is in conflict with the contextualist's view about the semantics of knowledge claims. For a helpful overview of this debate see Unger (1984, Ch. 1).

⁵⁰ In order for us to know that we are not dreaming, Stroud argues, we must be able to perform some test which determines that we are awake and not dreaming. However, in order for this test to be of any help at all, we would have to already know that we are not merely dreaming that we are performing the test, and/or dreaming that our test delivered us with the sought after result (Stroud 1984, 20-23). Recently, however, Stroud has backed off a bit from the claim that perception clearly cannot deliver us with the knowledge that we are not dreaming, and conceded, rightly to my mind, that it is the *assumption* of a certain view of perception which leads to this conclusion (see Stroud 2009a, 561-562). In this limited sense, then, Stroud is in fact in agreement with some of the substantive points I make below. However, as I hope to show, he fails to see the relevance that this admission has for drastically weakening the claim that skepticism is conditionally correct.

⁵¹ In a recent writing Stroud resists the reading which sees the skeptic as setting down "requirements" that any satisfactory philosophical explanation of our knowledge must meet. See Stroud (2009b, 590-592) for discussion. It seems that Stroud resists calling the relevant strictures "requirements" because he thinks the conditions are best seen as what we intuitively *already see* is necessary in order to satisfy our philosophical curiosity regarding the matter at hand. The conditions come *from within*, then, as opposed to being externally imposed as the term "requirements" may suggest. With this in mind, I will nevertheless continue to use the terms "requirements", "conditions", etc. in what follows.

⁵² See Williams (1996) for an in-depth discussion of this matter.

⁵³ For a recent defense of a contrary position which claims that the truth of certain external world claims is what best explains certain facts about our mental lives, see Vogel (2005, 1990). For a response see Fumerton (2005).

⁵⁴ The relevant different ways that one can beg the question are as follows. First, as already discussed, one can assume that one already has some piece of external world knowledge, and use that assumed knowledge to explain why one knows other things. Also, one can instead assume that one has some piece of basic entailing knowledge.

⁵⁵ See Pitcher (1971); Armstrong (1961).

⁵⁶ There are too many examples of this sort of view to exhaustively mention. Contemporary defenses of classical foundationalism in particular can be found in Fumerton (1995); Bonjour (2001a, 2001b, 2003a, 2003b); and McGrew et. al. (2007).

⁵⁷ The position sometimes (tellingly) referred to as *dogmatism*, however, maintains that perception alone can provide *prima facie justification* for certain external world beliefs. See Pryor (2000, 2004) and Huemer (2001) for a defense of this sort of position.

⁵⁸ See Stroud (1984, Ch. 3). For further discussion of what seems wrong with Moore's way with skepticism see Wright (1985, 2000, 2002, 2003, 2004); Pryor (2004); and Davies (1998, 2000, 2003, 2004).

⁵⁹ Namely, it has to be true that the only reason we would ever think exclusive reliance on our sense experiences was obviously not adequate for *establishing* that we are not

dreaming, is if we thought that some version of a conjunctive theory of perception was true. I argue for this claim below.

⁶⁰ See Pryor (2000); Huemer (2001).

⁶¹ For a discussion of whether experiences have contents, see Travis (2004); Siegel (2010b); and Byrne (2009).

⁶² Here I am employing Thompson Clarke's famous use of "plain"; see his (1972).

⁶³ Rorty (1979) famously discusses the historical role that a commitment to representationalist-style ideas has played in engendering perennial philosophical problems like the skeptical problem. The negative arguments given below are in roughly the same spirit. However, I think a powerful version of the skeptical problem still exists even after we accept a view like Rorty's. I discuss this matter in the final section of the paper.

⁶⁴ For an argument that arrives at a similar conclusion see Vinci (1986, 568-569).

⁶⁵ It is intelligible, after all, to entertain the story that right now I am in Vatican City talking to the Pope. But it simply doesn't follow that *were I to be in Vatican City talking to the Pope right now, then from a first-person standpoint I wouldn't know the difference*. In fact, the intuitive way to treat this last subjunctive conditional is to say that my sensory evidence would be quite different, were I involved in the scenario described in the antecedent. What this shows is that, generally speaking at least, the mere intelligibility of a described scenario at a certain time does not entail the reasonable acceptance of the subjunctive conditional involving that scenario in the antecedent obtaining at that time, and something like "from a first-person standpoint I wouldn't be able to tell the difference" or "My evidence, from a first-person standpoint, would be exactly the same" in the consequent.

⁶⁶ This sort of line of reasoning relates to Chisholm's discussion of perceptually-based justification in (1957, Ch. 5).

⁶⁷ For a similar point see Pitcher (1971, 101) and Williams (1996, 73-79). It should also be noted here that neither can the familiar Arguments from Illusion/Perceptual Relativity adequately support a conjunctive theory. This is because both cases of illusion (according to one common enough understanding of that term) as well as perceptually-relative experiences are compatible with a *physical existent* being the object of perception; hence both phenomena are obviously compatible with a non-cognitive theory. For a convincing argument for why neither of the above arguments work for establishing their desired conclusion, see Pitcher (1971, Ch. 1).

⁶⁸ For an argument for the same conclusion see Vinci (1986, 572-574).

⁶⁹ For arguments for a similar position see Lehrer (1971); Hilpinen (1983). Of course, those who endorse an appropriately understood KK thesis will claim that our inability to refute a second-order skeptic *entails* that we do not, in fact, know anything about external

reality. This is a paradigmatically internalist position. However, whether or not seeing skepticism as primarily a second-order affair requires *also* endorsing some kind of KK thesis is, I think, an open question. For the mere purposes of determining whether a constructive realist argument can be successful--which is, of course, at least one central way of determining the source of the skeptical problem--all that is important to realize is that the converse of a strong KK thesis obviously holds: knowing that I know that p entails knowing that p.

⁷⁰ See Wright (1985, 434-438); Williams (1996, 62-63). For a contrary view see Pryor (2000, 2004).

⁷¹ For helpful discussions of the so-called transparency of sense experience see Moore (1903); Harman (1990, 39); Tye (2000, 51-52); Dretske (1995, 62); and Kind (2003).

⁷² The original presentation of the new evil demon thought experiment can be found in Cohen and Lehrer (1983); see also Cohen (1984). For a well-known externalist response see Goldman (1986, 107; 1988).

⁷³ For discussion of cases of forgotten evidence, see Goldman (1999); Senor (1993) and Audi (1995).

⁷⁴ A prominent contemporary philosopher who has worked to develop this general kind of view of perception is Alva Noë. See, for example, Noë (2002a, 2002b, 2002c, 2004, and 2006). It is not clear to me, however, whether Noë would endorse a non-cognitive theory of perception in particular, as presented in this chapter. For general discussion of the connection between perceptual experience and action see also Ludwig (2006); O'Regan and Noë (2001); O'Shaughnessy (1992); and Viger (2006).

REFERENCES

- Adams, R. (1975) "Where Do Our Ideas Come From? — Descartes vs. Locke," in S. P. Stich, (ed.) *Innate Ideas* Berkeley: University of California Press.
- Alston, W. (1985) "Concepts of Epistemic Justification" in Crumley, J. (ed.) *Readings in Epistemology*, Mountain View, CA: Mayfield Publishing, 177 – 193.
- (1986) "Internalism and Externalism in Epistemology" in Kornblith, H. (ed.) *Epistemology: Internalism and Externalism* Malden, MA: Blackwell 2001, 68-110.
- (1988a) "An Internalist Externalism" *Synthese* 74, 265-283.
- (1988b) "Justification and Knowledge" *Proceedings of the World Congress of Philosophy* 5, Reprinted in Alston, (1989) 172-82.
- (1989) *Epistemic Justification* Ithaca: Cornell University Press.
- (1993) *The Reliability of Sense Perception* Ithaca, NY: Cornell University.
- (1995) "How to Think about Reliability" *Philosophical Topics* 23, 1-29.
- (2002) "Plantinga, Naturalism, and Defeat" in J. Beilby (ed.) *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument against Naturalism*. Ithaca: Cornell University Press, 176-203.
- (2005) *Beyond "Justification": Dimensions of Epistemic Evaluation* Ithaca, NY: Cornell University Press.
- Annis, D. (1973) "Knowledge and Defeasibility" *Philosophical Studies* 24, 199-203.
- Armstrong, D. M. (1961) *Perception and the Physical World* London: Routledge.
- (1968) *A Materialist Theory of the Mind* London: Routledge and Kegan Paul.
- (1973) *Belief, Truth and Knowledge* New York: Cambridge University Press.
- Audi, R. (1995) "Memorial Justification" *Philosophical Topics* 23, 31-45.
- Austin, J. L. (1961) 'Other Minds' in his *Philosophical Papers*, Oxford: Oxford University Press.
- (1962) *Sense and Sensibilia* Oxford: Oxford University Press.

- Ayer, A. J., (1946-7) "Phenomenalism," *Proceedings of the Aristotelian Society* 47, 163-96.
- (1956) *The Problem of Knowledge*, London: Cambridge University Press.
- Bach, K. (1985) "A Rationale for Reliabilism." *The Monist* 68, 246-263.
- Barker, J. (1976) "What You Don't Know Won't Hurt You" *American Philosophical Quarterly* 13, 303-308.
- Bennett, J. (1990) "Truth and Stability in Descartes' *Meditations*," *Canadian Journal of Philosophy*, 16 (supplement): 75-108.
- Bergmann, M. (1997a) "Internalism, Externalism, and Epistemic Defeat." (PhD Dissertation: University of Notre Dame).
- (1997b) "Internalism, Externalism, and the No-Defeater Condition" *Synthese* 110, 399-417.
- (2000) "Deontology and Defeat" *Philosophy and Phenomenological Research* 60, 87-102.
- (2005) "Defeaters and Higher-Level Requirements" *Philosophical Quarterly* D5, 419-36.
- (2006) *Justification without Awareness*, New York: Oxford University Press.
- Beysade, M. (1993) "Privileged Truth or Exemplary Truth?" in S. Voss (ed.) *Essays on the Philosophy and Science of René Descartes*, Oxford: Oxford University Press.
- Black, T. and P. Murphy (2007) "In Defense of Sensitivity" *Synthese* 154, 53-71.
- Block, N. (1990) "Inverted Earth" in J. Tomberlin (ed.) *Philosophical Perspectives* 4, 53-79.
- (1996) "Mental Paint and Mental Latex" in E. Villanueva (ed.) *Philosophical Issues* 7 Atascadero, CA: Ridgeview (19-49).
- (1997) "Inverted Earth" in N. Block, O. Flanagan and G. Güzeldere (eds.) *The Nature of Consciousness* Cambridge, MA: MIT Press.
- (2004) "Mental Paint", in M. Hahn and B. Ramberg (ed.) *Others on Burge: 10 Essays With Responses From Tyler Burge*, Cambridge, MA: MIT Press.

- BonJour, L. (1980) "Externalists Theories of Empirical Knowledge" in Kornblith, H. (ed.) *Epistemology: Internalism and Externalism*, Malden, MA: Blackwell (2001), 10-35.
- (1985) *The Structure of Empirical Knowledge* Cambridge, MA: Harvard University Press.
- (2001a) 'Toward a Defense of Empirical Foundationalism' in DePaul, M. (ed.) *Resurrecting Old-Fashioned Foundationalism*, Lanham, MD: Rowman & Littlefield, 21-40.
- (2001b) 'Replies to Pollock and Plantinga' in *Resurrecting Old-Fashioned Foundationalism*, 79-86.
- (2003a) 'A Version of Internalist Foundationalism' in Bonjour, L. and E. Sosa (eds.) *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, Malden, MA: Blackwell, 3-96.
- (2003b) 'Reply to Sosa' in *Epistemic Justification: Internalism vs. Externalism*, 173-200.
- (2004) "In Search of Direct Realism," *Philosophy and Phenomenological Research*, 69, 349-67.
- BonJour, L., and E. Sosa (2003). *Epistemic Justification* Malden, MA: Blackwell.
- Bouwsma, O. K. (1949) "Descartes' Evil Genius" *Philosophical Review* 58, 141-151.
- Brueckner, A. (1992) "Problems with the Wright Route to Skepticism" *Mind* 101, 309-17.
- Brewer, B. (2000) *Perception and Reason* Oxford: Oxford University Press.
- Broughton, J. (2002) *Descartes's Method of Doubt*, Princeton University Press.
- Burge, T. (1982) "Other Bodies" in A. Woodfield (ed.) *Thought and Object* Oxford: Clarendon Press.
- (1986) "Cartesian Error and the Objectivity of Perception" in P. Pettit and J. McDowell (eds.) *Subject, Thought and Content* Oxford University Press, 117-36.
- (1991) "Vision and Intentional Content" in LePore, E. and R. Van Gulick (eds.) *John Searle and his Critics* Oxford: Blackwell.
- Byrne, A. (2004) "How Hard are the Sceptical Paradoxes?" *Nous* 38, 299-325.

- (2009) "Experience and Content" *Philosophical Quarterly* 59, 429-51.
- Carroll, L. (1985) "What the Tortoise Said to Achilles" *Mind* 4, 278-280.
- Carriero, J. (2009) *Between Two Worlds: A Reading of Descartes' Meditations*, Princeton University Press.
- Cavell, S. (1979) *The Claim of Reason* Oxford: Oxford University Press.
- Chalmers, D. (2006) "Perception and the Fall from Eden" in Gendler, T. and J. Hawthorne (eds.) *Perceptual Experience*, Oxford: Clarendon Press, 49-125.
- Chappell, Vere (1986) "The Theory of Ideas" in A. O. Rorty (ed.) *Essays on Descartes' Meditations* Berkeley: University of California Press.
- Chisholm, R. (1942) "The Problem of the Speckled Hand" *Mind* 51, 368-373.
- (1957) *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- (1977) *Theory of Knowledge* 2nd edition, Englewood Cliffs: Prentice-Hall.
- (1982) *The Foundations of Knowing*, Minneapolis: University of Minnesota Press.
- (1988) "The Indispensability of Internal Justification" *Synthese* 74, 285-296.
- Clarke, T. (1972) "The Legacy of Skepticism" *Journal of Philosophy* 69, 754-69.
- Coffman, E. J. (2007) "Thinking about Luck" *Synthese* 158, 385-398.
- Cohen, S. (1984) "Justification and Truth" *Philosophical Studies* 46, 279-295.
- (1988) "How to be a Fallibilist" *Philosophical Perspectives* 2, 91-123.
- (1998) "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery" *Australasian Journal of Philosophy* 76, 289-306.
- (1999) "Contextualism, Skepticism, and the Structure of Reasons" *Philosophical Perspectives* 13, 57-89.
- Conee, E. and R. Feldman (1985) "Evidentialism" *Philosophical Studies* 48, 15-34.
- (2004a) *Evidentialism: Essays in Epistemology* New York: Oxford University Press.
- (2004b) "Internalism Defended" in *Evidentialism: Essays in Epistemology*, 53-82.

- Crane, Tim (1992) *The Contents of Experience* (ed.) Cambridge: Cambridge University Press.
- (1998) "Content, Non-Conceptual" in E.J. Craig (ed.) *Routledge Encyclopedia of Philosophy* London: Routledge.
- (2001) *Elements of Mind* Oxford: Oxford University Press.
- Crane, T. and K. Farkas (2004) *Metaphysics: a Guide and Anthology* (eds.) Oxford: Oxford University Press.
- Cunning, D. (2007) "Descartes on the Dubitability of the Existence of Self," *Philosophy & Phenomenological Research* 74, 111–131.
- Curley, E. M. (1978) *Descartes Against the Sceptics*, Cambridge, MA: Harvard University Press.
- (1986) "Analysis in the *Meditations*: The Quest for Clear and Distinct Ideas," in A. O. Rorty (ed.) *Essays on Descartes' Meditations*, Berkeley: University of California Press.
- (1993) "Certainty: Psychological, Moral, and Metaphysical," in S. Voss (ed.) *Essays on the Philosophy and Science of René Descartes*, Oxford: Oxford University Press.
- Dancy, J. (1988) *Perceptual Knowledge* (ed.) Oxford: Oxford University Press.
- (1995) "Arguments from Illusion" *The Philosophical Quarterly* 45, 421–38.
- Davidson, D. (1989) "A Coherence Theory of Truth and Knowledge" in LePore, E. (ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, New York: Blackwell, 307-319.
- Davies, M. (1992) "Perceptual Content and Local Supervenience" *proceedings of the Aristotelian Society* 92, 21–45.
- (1998) 'Externalism, Architecturalism, and the Epistemic Warrant' in Wright, C., M. Smith and C. Macdonald (eds.) *Knowing Our Own Minds: Essays in Self-Knowledge*, Oxford University Press, 321-361.
- (2000) "Externalism and Armchair Knowledge' in Boghossian, P., and C. Peacocke (eds.) *New Essays on the A Priori*, Oxford University Press, 384-414.

- (2003) 'The Problem of Armchair Knowledge' in Nuccetelli, S. (ed.) *New Essays on Semantic Externalism and Self-Knowledge*, MIT Press, 23-55.
- (2004) 'Epistemic Entitlement, Warrant Transmission, and Easy-Knowledge' *Aristotelian Society Supplement* 78, 213-45.
- Della Rocca, M. (2005) "Descartes, the Cartesian Circle, and Epistemology Without God," *Philosophy and Phenomenological Research* 70, 1-33.
- DeRose, K. (1992a) "Descartes, Epistemic Principles, Epistemic Circularity, and Scientia" *Pacific Philosophical Quarterly* 73, 220-38.
- (1992b) "Contextualism and Knowledge Attributions" *Philosophy and Phenomenological Research* 52, 913-929.
- (1995) "Solving the Skeptical Problem" in DeRose, K. and T. Warfield (eds.) *Skepticism: A Contemporary Reader*, Oxford: Oxford University Press (1999), 183-219.
- (2002) "Assertion, Knowledge and Context" *Philosophical Review* 111, 167-203.
- Descartes, R. (1641) *Meditations on First Philosophy*, in J. Cottingham, R. Stoothoff, and D. Murdoch (trans.) *The Philosophical Writings of Descartes vol. II* New York: Cambridge University Press (1984, 1-62.
- Dewey, J. (1917) "The Need for a Recovery of Philosophy" in McDermott, J. (ed.) *The Philosophy of John Dewey: Vols. I and II* Chicago and London: The University of Chicago Press, 58-97.
- (1929) *The Quest for Certainty: A Study of the Relation of Knowledge and Action*, New York: Capricorn Books.
- Doney, W. (1955) "The Cartesian Circle" *Journal of the History of Ideas* 16, 324-38.
- (1987) *Eternal Truths and the Cartesian Circle*, New York: Garland Publishing.
- Dretske, F. (1969) *Seeing and Knowing* Chicago: University of Chicago Press.
- (1970) "Epistemic Operators" *Journal of Philosophy* 67, 1007-1023.
- (1981) *Knowledge and the Flow of Information* Oxford: Blackwell.
- (1995) *Naturalizing the Mind* Cambridge, Mass.: MIT Press.

- (2005) "The Case Against Closure" in Steup, M. and E. Sosa (eds.) *Contemporary Debates in Epistemology*, Malden, MA: Blackwell, 13-25.
- Ducasse, C.J. (1942) "Moore's Refutation of Idealism" in Schilpp, P. (ed.) *The Philosophy of G.E. Moore* Chicago: Northwestern University Press.
- Dunlop, C. E. M. (1977) *Philosophical Essays on Dreaming*, Ithaca: Cornell University Press.
- Fales, E. (1996) *A Defense of the Given*, Lanham MD: Rowman & Littlefield.
- Feldman, R. (1981) "Fallibilism and Knowing That One Knows" *The Philosophical Review* 90, 266-282.
- (2002) *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- (2004a) "In Search of Internalism and Externalism" in R. Schantz (ed.) *The Externalist Challenge* New York: Walter de Gruyter, 143-156.
- (2004b) "Having Evidence" in Conee and Feldman (eds.) *Evidentialism: Essays in Epistemology*, 219-241.
- (2004c) "Foundational Beliefs and Empirical Possibilities" *Philosophical Issues* 14, 132-147.
- (2005a) "Respecting the Evidence" *Philosophical Perspectives* (19, 95-119).
- (2005b) "Justification is Internal" in M. Steup and E. Sosa (eds.) *Contemporary Debates in Epistemology* Malden, MA: Blackwell, 270-284.
- Fetzer, J. H. (1982) "Probabilistic Explanations", *PSA: Proceedings of the Biennial Meeting of Philosophy of Science Association* 2, 194-207.
- (1983) "Probability and Objectivity in Deterministic and Indeterministic Situations" *Synthese* 57, 367-386.
- Foley, R. (1993) "What Am I to Believe?" in S. Wagner and R. Warner (eds.) *Naturalism: A Critical Appraisal* University of Notre Dame Press, 147-162.
- Foster, J. (2000) *The Nature of Perception* Oxford: Oxford University Press.
- Frankfurt, H. G. (1970) *Demons, Dreamers, and Madmen*, Indianapolis: The Bobbs-Merrill Company.

- (1997) "Descartes on the Real Existence of Matter," *Topoi* 16, 153–162.
- Fumerton, R. (1985) *Metaphysical and Epistemological Problems of Perception* Lincoln: University of Nebraska Press.
- (1988) "The Internalism/Externalism Controversy" *Philosophical Perspectives* 2, 443-459.
- (1995) *Metaepistemology and Skepticism* Lanham, MD: Rowman and Littlefield.
- (2001) "Replies to Pollock and Plantinga" in Depaul, M. (ed.) *Resurrecting Old-Fashioned Foundationalism*, Lanham, MD: Rowman and Littlefield, 69-78.
- (2004) "Inferential Internalism and the Presuppositions of Skeptical Arguments" in R. Schantz (ed.) *The Externalist Challenge*, 157-167.
- (2005) 'The Challenge of Refuting Skepticism' in Steup, M. and E. Sosa (eds.) *Contemporary Debates in Epistemology*, Malden, MA: Blackwell Publishing, 85-97.
- (2006) "Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament" in T. Crisp, M. Davidson, and D. vander Laan (eds.) *Knowledge and Reality: Essays in honor of Alvin Plantinga, Philosophical Studies Series* 103, Dordrecht: Springer, 179-191.
- Garber, D. (1986) "*Semel in vita*: The Scientific Background to Descartes' *Meditations*," in A. O. Rorty (ed.) *Essays on Descartes' Meditations*, Berkeley: University of California Press.
- (1992) *Descartes' Metaphysical Physics*, Chicago: University of Chicago Press.
- Gaukroger, S. (1989) *Cartesian Logic: An Essay on Descartes's Conception of Inference*, Oxford: Clarendon Press.
- Giere, R. N. (1973) "Objective Single-Case Probabilities and the Foundations of Statistics" in P. Suppes *et al.*, (eds.) *Logic, Methodology and Philosophy of Science* (Volume IV), New York: North-Holland.
- Gettier, E. (1963) "Is Justified True Belief Knowledge?" *Analysis* 23, 121-123.
- Gewirth, A. (1941) "The Cartesian Circle," *Philosophical Review* 50, 368–95.
- Gillies, D. (2000) "Varieties of Propensity" *British Journal for the Philosophy of Science* 51, 807–835.

- Ginet, C. (1975) *Knowledge, Perception, and Memory* Dordrecht: D. Reidel.
- Goldberg, S. (2007) "How Lucky Can You Get?" *Synthese* 158, 313-327.
- Goldman, A. (1967) "A Causal Theory of Knowing" *The Journal of Philosophy* 64, 357-372.
- (1976) "Discrimination and Perceptual Knowledge" *Journal of Philosophy* 73, 771-91.
- (1979) "What is Justified Belief?" in G. S. Pappas (ed.) *Justification and Knowledge* Dordrecht: D. Reidel, 1-23.
- (1980) "The Internalist Conception of Justification" in Kornblith, H. (ed.) *Epistemology: Internalism and Externalism*, Malden, MA: Blackwell (2001), 36-67.
- (1986) *Epistemology and Cognition* Cambridge, MA: Harvard University Press.
- (1988) "Strong and Weak Justification." *Philosophical Perspectives* 2, 51-69.
- (1993) "Epistemic Folkways in Scientific Epistemology" *Philosophical Issues* 3, 271-285.
- (1999) "Internalism Exposed" *Journal of Philosophy* 96, 271-293.
- Greco, J. (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry* Cambridge: Cambridge University Press.
- (2003) "Virtue and Luck, Epistemic and Otherwise" *Metaphilosophy* 34, 353-366.
- (2007) "Worries about Pritchard's Safety" *Synthese* 158, 299-302.
- Grice, H.P. (1961) "The Causal Theory of Perception" *Proceedings of the Aristotelian Society Supplementary Volume* 35: 121-52.
- Hacking, I. (1965) *The Logic of Statistical Inference*, Cambridge: Cambridge University Press.
- (1980) "Proof and Eternal Truths: Descartes and Leibniz," in S. Gaukroger (ed.) *Descartes: Philosophy, Mathematics and Physics*, Sussex: The Harvester Press.
- Hambourger, R. (1987) "Justified Assertion and the Relativity of Knowledge" *Philosophical Studies* 51, 241-269.
- Harman, G. (1973) *Thought*, Princeton, NJ: Princeton University Press.

- (1990) 'The Intrinsic Quality of Experience' in Tomberlin, J. (ed.) *Philosophical Perspectives* 4, Atascadero, CA: Ridgeview Press.
- Hawthorne, J. (2004) *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Hetherington, S. (1990) "Epistemic Internalism's Dilemma" *American Philosophical Quarterly* 27, 245-251.
- (1991) "On Being Epistemically Internal" *Philosophy and Phenomenological Research* 51, 855-871.
- Hilbert, D. "Hallucination, Sense-Data and Direct Realism" *Philosophical Studies* 120, 185-191.
- Hiller, A., and Neta, R. (2007) "Safety and Epistemic Luck" *Synthese* 158, 303-313.
- Hilpinen, R. (1983) 'Skepticism and Justification' *Synthese* 55, 165-173.
- Hintikka, J. (1962a) "Cogito ergo sum: Inference or Performance?" *Philosophical Review* 71, 3-32.
- (1978) "A Discourse on Descartes's Method," in M. Hooker (ed.) *Descartes: Critical and Interpretive Essays*, Baltimore: Johns Hopkins University Press.
- Hinton, J.M. (1973) *Experiences* Oxford: Clarendon Press.
- Hoffman, P. (1996) "Descartes on Misrepresentation," *Journal of the History of Philosophy* 34, 357-381.
- Huemer, M. (2001) *Skepticism and the Veil of Perception* Lanham, MD: Rowman & Littlefield.
- (2002) "Fumerton's Principle of Inferential Justification" *Journal of Philosophical Research* 27, 329-40.
- Hume, D. (1748) *An Enquiry Concerning Human Understanding*, T. Beauchamp (ed.) Oxford: Oxford University Press (1999).
- Jackson, F. (1977) *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Johnston, M. (2004) "The Obscure Object of Hallucination" *Philosophical Studies* 103:113-83.

- Jolley, N. (1990) *the Light of the Soul: Theories of Ideas in Leibniz, Malebranche, and Descartes*, Oxford: Clarendon Press.
- Kenny, A. (1968) *Descartes: A Study of His Philosophy*, New York: Random House.
- Kind, A. (2003) 'What's So Transparent about Transparency?' *Philosophical Studies* 115, 225-244.
- Klein, P. (1971) "A Proposed Definition of Propositional Knowledge" *The Journal of Philosophy* 68, 471-82.
- (1976) "Knowledge, Causality, and Defeasibility" *The Journal of Philosophy* 73, 792-812.
- (1981) *Certainty: A Refutation of Skepticism* Minnesota: University of Minnesota Press.
- (1998) "Foundationalism and the Infinite Regress of Reasons" *Philosophy and Phenomenological Research* 58, 19-26.
- Korcz, K. (2000) "The Causal-Doxastic Theory of the Basing Relation" *Canadian Journal of Philosophy* 30, 525 – 550.
- Kornblith, H. (1988) "How Internal Can You Get?" *Synthese* 74, 313-327.
- (2001) *Epistemology: Internalism and Externalism* (ed.) Malden, MA: Blackwell.
- Kvanvig, J. (2003) "Propositionalism and the Perspectival Character of Justification" *American Philosophical Quarterly* 40, 3-18.
- Lehrer, K. (1971) 'Why Not Scepticism?' *Philosophical Forum* 2, 283-298.
- (1974) *Knowledge*, Oxford: Clarendon Press.
- Lehrer, K. and S. Cohen (1983) "Justification, Truth and Knowledge" *Synthese* 55 (191-207).
- Lehrer, K. and Paxson, T. (1969) "Knowledge: Undefeated Justified True Belief" *Journal of Philosophy* 66, 225-37.
- Leite, A. (2004) 'Is Fallibility an Epistemological Shortcoming?' *The Philosophical Quarterly* 54, 232-251.

- Lennon, T. M. (2008) *The Plain Truth: Descartes, Huet, and Skepticism*, Leiden: Brill.
- Lewis, D. (1973) *Counterfactuals*, Cambridge, MA: Harvard University Press.
- (1996) "Elusive Knowledge" *Australasian Journal of Philosophy* 74, 549-567.
- Loar, B. (2002) "Transparent Experience" in A. Jokic and Q. Smith (eds.) *Consciousness: New Philosophical Perspectives* Oxford: Oxford University Press.
- Loeb, L. E. (1992) "The Cartesian Circle," in J. Cottingham (ed.) *The Cambridge Companion to Descartes*, Cambridge: Cambridge University Press.
- Ludwig, K. (2006) "Is the Aim of Perception to Provide Accurate Representations?" in R. Stainton (ed.) *Contemporary Debates in Cognitive Science*, Malden, MA: Blackwell Publishing, 259-274.
- Lycan, W.G. (1996) *Consciousness and Experience* Cambridge, MA: MIT Press.
- (1999) *Mind and Cognition* Second Ed. Oxford: Blackwell.
- Markie, P. (1992) "The Cogito and Its Importance" in J. Cottingham (ed.) *The Cambridge Companion to Descartes*, Cambridge: Cambridge University Press.
- Martin, M.G.F. (1992a) "Sight and Touch" in Crane (ed.) (1992).
- (1992b) "Perception, Concepts and Memory" *Philosophical Review* 101, 745-63.
- (1998) "Setting Things Before the Mind" in A. O'Hear (ed.) *Contemporary Issues in the Philosophy of Mind* Cambridge: Cambridge University Press.
- (2002) "The Transparency of Experience" *Mind and Language* 17, 376-425.
- McCulloch, Gregory (1995, *The Mind and its World* London: Routledge.
- (2003) *The Life of the Mind* London: Routledge.
- McDermid, D. J. (1994) "The Content of Perceptual Experience" *Philosophical Quarterly* 44 (190-205; reprinted in Noë & Thompson (eds.) 2002.
- (1994a) *Mind and World* Cambridge, MA: Harvard University Press.
- McGinn, C. (1989) *Mental Content* Oxford: Blackwell.
- McGinn, M. (1989) *Sense and Certainty* Oxford: Basil Blackwell.

- McGrew, T. and L. McGrew (2007) *Internalism and Epistemology: The Architecture of Reason*, New York: Routledge.
- Menn, S. (1998) *Descartes and Augustine*, Cambridge University Press.
- Merleau-Ponty, M. (1945) *The Phenomenology of Perception*, Colin Smith (trans.), London: Routledge 2002.
- Miller, D. W. (1994) *Critical Rationalism: A Restatement and Defence*, Lasalle, IL: Open Court.
- Millikan, R. G. (1987) *Language, Thought and Other Biological Categories* Cambridge, MA: MIT Press.
- Moore, G. E. (1903) 'The Refutation of Idealism' in *Philosophical Studies*, London: Routledge.
- (1939) 'Proof of an External World' *Philosophical Papers*, London: Allen & Unwin, 127-150.
- Morris, J. (1973) "Descartes' Natural Light" *Journal of the History of Philosophy* 11, 169-187.
- Moser, P. (1989) *Knowledge and Evidence*, Cambridge: Cambridge University Press.
- Moyal-Sharrock, D. (2000) "Wittgenstein Distinguished: A Response to Pieranna Garavaso on Wittgenstein and Quine" *Philosophical Investigations* 23, 54-69.
- (2003) "Logic in Action: Wittgenstein's Logical Pragmatism and the Impotence of Scepticism" *Philosophical Investigations* 26, 125-148.
- (2004) *The Third Wittgenstein: The Post-Investigations Works*, London: Ashgate.
- (2007) *Understanding Wittgenstein's On Certainty*, London: Palgrave Macmillan.
- Moyal-Sharrock, D. and W. Brenner (2005) *Readings of Wittgenstein's On Certainty*, New York: Palgrave Macmillan.
- Nadler, S. (2006) "The Doctrine of Ideas," in S. Gaukroger (ed.) *The Blackwell Guide to Descartes' Meditations*, Blackwell Publishing.
- Nagel, T. (1986) *The View from Nowhere* Oxford: Oxford University Press.

- Nelson, A. (1997) "Descartes's Ontology of Thought" *Topoi* 16, 163–178.
- Neta, R. (2003) "Contextualism and the Problem of the External World" *Philosophy and Phenomenological Research* 66, 1-31.
- Newman, L. (1994) "Descartes on Unknown Faculties and Our Knowledge of the External World" *Philosophical Review* 103, 489–531.
- (1999) "The Fourth Meditation" *Philosophy & Phenomenological Research* 59, 559–591.
- (2004) "Rocking the Foundations of Cartesian Knowledge: Critical Notice of Janet Broughton, *Descartes's Method of Doubt*" *Philosophical Review* 113, 101–125.
- (2006) "Descartes' Rationalist Epistemology," in A. Nelson (ed.) *A Companion to Rationalism*, Blackwell Companions to Philosophy.
- (2007) "Descartes on the Will in Judgment," in J. Broughton and J. Carriero (eds.) *A Companion to Descartes* Blackwell Companions to Philosophy.
- (2009) "Ideas, Pictures, and the Directness of Perception in Descartes and Knowledge" *Philosophy Compass* 4, 134–54.
- Newman, L. and A. Nelson (1999) "Circumventing Cartesian Circles" *Noûs* 33, 370–404.
- Noë, A. (2002a) "Is the Visual World a Grand Illusion?" *Journal of Consciousness Studies* 9, 1-12.
- (2002b) *Is the Visual World a Grand Illusion?* (ed.) Special issue of *Journal of Consciousness Studies* 9, Charlottesville, VA: Imprint Academic.
- (2002c) "On What We See" *Pacific Philosophical Quarterly* 83, 57-80.
- (2004) *Action in Perception*, Cambridge, MA: MIT Press.
- (2006) "Experience without the Head" in T. S. Gendler and J. Hawthorne (eds.) *Perceptual Experience*, Oxford: Oxford University Press, 411-433.
- Noë, A. and O'Regan, J. K. (2001) "A Sensorimotor Approach to Vision and Visual Consciousness" *Behavioral and Brain Sciences* 24, 883-975.
- Noë, A. and E. Thompson (eds.) (2002, *Vision and Mind: Selected Readings in the Philosophy of Perception* Cambridge, MA: MIT Press.

- Nolan, L. and A. Nelson (2006) "Proofs for the Existence of God" in S. Gaukroger (ed.) *The Blackwell Guide to Descartes' Meditations*, Oxford: Blackwell Publishing.
- Nozick, R. (1981) *Philosophical Explanations* Cambridge, MA: Belknap Press.
- O'Hara, K. (1993) "Sceptical Overkill: On Two Recent Arguments Against Scepticism" *Mind* New Series 102, 315-327.
- O'Shaughnessy, B. (1992) "The Diversity and Unity of Action and Perception" in T. Crane (ed.) *The Contents of Experience: Essays on Perception*, Cambridge: Cambridge University Press, 216-266.
- Pappas, G. (1979) "Basing Relations" in Pappas (ed.) *Justification and Knowledge*, Dordrecht: D. Reidel Publishing Co.
- Pappas, N. (1991) 'Knowing and Saying I Know' *Philosophy* 66, 487-502.
- Parrini, P., W. Salmon, and M. Salmon, (eds.) (2003) *Logical Empiricism: Historical and Contemporary Perspectives*, Pittsburgh: University of Pittsburgh Press.
- Pautz, A. (2010) 'Why Explain Visual Experience in Terms of Content?' In B. Nanay (ed.) *Perceiving the World*, New York: Oxford University Press, 254-309.
- Peacocke, C. (1983) *Sense and Content* Oxford: Oxford University Press.
- (1992) *A Study of Concepts* Cambridge, MA: MIT Press.
- Perkins, M. (1983) *Sensing the World* Indianapolis: Hackett.
- Pitcher, G. (1970) *A Theory of Perception* Princeton: Princeton University Press.
- Plantinga, A. (1993) *Warrant: The Current Debate* New York: Oxford University Press.
- Pollock, J. (1970) "The Structure of Epistemic Justification" *American Philosophical Quarterly*, monograph series 4, 62-78.
- (1974) *Knowledge and Justification*, Princeton, NJ: Princeton University Press.
- (1984) "Reliability and Justified Belief" *Canadian Journal of Philosophy* 14, 103-114.
- (1986) *Contemporary Theories of Knowledge*, Savage, MD: Rowman and Littlefield.

- Pollock, J. and J. Cruz (1999) *Contemporary Theories of Knowledge*, 2nd ed. New York: Rowman & Littlefield.
- Popkin, R. H. (1979) *The History of Scepticism from Erasmus to Spinoza*, Berkeley: University of California Press.
- Popper, K. R. (1957) "The Propensity Interpretation of the Calculus of Probability and the Quantum Theory", in S. Körner (ed.), *The Colston Papers* 9, 65–70.
- Price, H.H. (1932) *Perception*, London: Methuen.
- Prichard, H.A. (1950) *Knowledge and Perception* Oxford: Clarendon Press.
- Pritchard, D. (2001) "Radical Scepticism, Epistemological Externalism, and "Hinge" Propositions" *Wittgenstein-Jahrbuch 2001/2002*, (ed.) D. Salehi,
- (2003) "Virtue Epistemology and Epistemic Luck" *Metaphilosophy* 34,106-130.
- (2005a) *Epistemic Luck*, Oxford: Oxford University Press.
- (2005b) "Scepticism, Epistemic Luck and Epistemic Angst" *Australasian Journal of Philosophy* 83, 185-206.
- (2005c) "Wittgenstein's On Certainty and Contemporary Anti-Scepticism" in D. Moyal-Sharrock & W. H. Brenner (eds.) *Investigating On Certainty: Essays on Wittgenstein's Last Work*, New York: Palgrave Macmillan.
- (2007) "Anti-Luck Epistemology" *Synthese* 158, 277-297.
- Pryor, J. (2000) "The Skeptic and the Dogmatist," *Noûs* 34, 517-49.
- (2004) 'What's Wrong with Moore's Argument?' *Philosophical Issues* 14, 349-378.
- Quine, W. V. O. (1960) *Word and Object*, Cambridge: MIT Press.
- (1969) *Ontological Relativity and other Essays*, New York: Columbia University Press.
- (1974) *The Roots of Reference*, La Salle, IL: Open Court Publishing.
- (1975) "The Nature of Natural Knowledge" in S. Guttenplan (ed.) *Mind and Language*, Oxford: Clarendon Press, 67-81.

- (1976) *The Ways of Paradox and other Essays*, Cambridge: Harvard University Press.
- (1980) *From a Logical Point of View*, Cambridge: Harvard University Press.
- (1981a) *Theories and Things*, Cambridge: Belknap Press.
- (1981b) "Reply to Stroud" *Midwest Studies in Philosophy* 6, 473-476.
- (1982) *Methods of Logic*, Cambridge: Harvard University Press.
- Reichenbach, H. (1949) *The Theory of Probability*, Berkeley: University of California Press.
- Rickless, S. C. (2005) "The Cartesian Fallacy Fallacy" *Noûs* 39, 309–336.
- Rorty, R. (1979) *Philosophy and the Mirror of Nature*, Princeton, NJ: Princeton University Press.
- Rott, H. (2004) "Stability, Strength and Sensitivity: Converting Belief into Knowledge" *Erkenntnis* 61, 469-493.
- Rudd, A. (2005) "Wittgenstein, Global Scepticism and the Primacy of Practice" in D. Moyal-Sharrock and W. Brenner (eds.) *Readings of Wittgenstein's On Certainty*, London and Basingstoke: Palgrave, 142-161.
- Russell, B. (1910-11) "Knowledge by Acquaintance and Knowledge by Description." *The Proceedings of the Aristotelian Society* 11, 209-32.
- (1912) *The Problems of Philosophy*, Oxford: Oxford University Press.
- (1945) *A History of Western Philosophy*, New York: Simon and Schuster.
- (1948) *Human Knowledge: Its Scope and Limits*, New York: Simon and Schuster.
- (1984) *Theory of Knowledge: The 1913 Manuscript*, E. Eames (ed.) London: Allen and Unwin Ltd.
- Ryle, G. (1949) *The Concept of Mind*, Chicago: The University of Chicago Press.
- Rysiew, P. (2001) "The Context-Sensitivity of Knowledge Attributions" *Noûs* 35, 477-514.

- Sarkar, H. (2003) *Descartes' Cogito: Saved from the Great Shipwreck*, Cambridge University Press.
- Searle, J. (1983) *Intentionality* Cambridge University Press.
- Sellars, W. (1963) *Science, Perception, and Reality*. London: Routledge and Kegan Paul.
- Senor, T. (1993) "Internalism Foundationalism and Justification of Memory Belief" *Synthese* 94, 453-76.
- Schiffer, S. (1996) "Contextualist Solutions to Scepticism" *Proceedings of the Aristotelian Society* 96, 317-333.
- Shoemaker, S. (1990) "Qualities and Qualia: What's in the Mind?" *Philosophy and Phenomenological Research* Suppl. Vol. 50, 109-131.
- (1994) "Self-Knowledge and 'Inner Sense'" *Philosophy and Phenomenological Research* 54, 249- 314.
- Shope, R. (1983) *The Analysis of Knowing: A Decade of Research*, Princeton, NJ: Princeton University Press.
- Siegel, S. (2010a) "The Contents of Perception " *The Stanford Encyclopedia of Philosophy*, E. Zalta (ed.) URL = <http://plato.stanford.edu/archives/fall2010/entries/perception-contents/>.
- (2010b) 'Do Experiences Have Contents?' in B. Nanay (ed.) *Perceiving the World*, New York: Oxford University Press, 333-368.
- (2010c) *The Contents of Visual Experience*, New York: Oxford University Press.
- Smith, A.D. (2002) *The Problem of Perception* Cambridge, Mass.: Harvard University Press.
- Snowdon, P.F. (1990) "The Objects of Perceptual Experience" *Proceedings of the Aristotelian Society Supplementary Volume* 64: 121–150.
- (1992) "How to Interpret 'Direct Perception'" in Crane (ed.) (1992).
- Sosa, E. (1980) "The Raft and the Pyramid" *Midwest Studies in Philosophy* 5, 3-25.
- (1991a) "Reliabilism and Intellectual Virtue" in *Knowledge in Perspective: Selected Essays in Epistemology* New York: Cambridge University Press, 131-145.

- (1991b) "Knowledge and Intellectual Virtue" in *Knowledge in Perspective*, 225-244.
- (1993) "Skepticism and Our Knowledge Circle" *Philosophical Issues* 4, 211-226.
- (1994) "Philosophical Scepticism and Epistemic Circularity" in K. DeRose and T. Warfield (eds.) *Skepticism: A Contemporary Reader*, Oxford: Oxford University Press (1999), 93-114.
- (1997a). "How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes," *Philosophical Studies* 85, 229-49.
- (1997b) "Reflective Knowledge in the Best Circles" *Journal of Philosophy* 94, 410-430.
- Spener, M. (2003) *Gilding or Staining the Mind: Phenomenology and the Metaphysics of Visual Experience* PhD thesis, University of London.
- Stanley, J. (2005) *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Steup, M. (1989) "The Regress of Metajustification" *Philosophical Studies* 55, 41-56.
- (1996) *An Introduction to Contemporary Epistemology*, Upper Saddle River, NJ: Prentice Hall.
- (1999) "A Defense of Internalism" in *The Theory of Knowledge: Classical and Contemporary Readings* 2nd edition, Belmont, CA: Wadsworth, 373-384.
- (2001) "Epistemic Duty, Evidence, and Internality" in M. Steup (ed.) *Knowledge, Truth, and Duty* New York: Oxford University Press.
- Strawson, P.F. (1979) "Perception and its Objects" in G. Macdonald (ed.) *Perception and Identity: Essays Presented to A.J. Ayer with His Replies* London: Macmillan.
- Stroll, A. (1994) *Moore and Wittgenstein on Certainty*, New York: Oxford University Press.
- Stroud, B. (1984) *The Significance of Philosophical Scepticism* Oxford: Clarendon Press.
- (1989) "Understanding Human Knowledge in General" in M. Clay and K. Lehrer (eds.) *Knowledge in Skepticism* Boulder: Westview Press.
- (1994) "Scepticism, 'Externalism', and the Goal of Epistemology" *Proceedings of the Aristotelian Society* 68, 291-307.

- (2000) *Understanding Human Knowledge* Oxford: Oxford University Press.
- (2009a) "Skepticism and the Senses" *European Journal of Philosophy* 17, 559-570.
- (2009b) 'Explaining Perceptual Knowledge: Reply to Quassim Cassam' *European Journal of Philosophy* 17, 590-596.
- Sturgeon, S. (1998) "Visual Experience" *Proceedings of the Aristotelian Society* 98: 179–200.
- Swain, M. (1974) "Epistemic Defeasibility" *The American Philosophical Quarterly* 11, 15-25.
- (1979) "Justification and the Basis of Belief" in Pappas, G. (ed.) *Justification and Knowledge*, Boston: D. Reidel Publishing Co.
- (1981) *Reasons and Knowledge* Ithaca: Cornell University Press.
- Swartz, R.J. (1965) *Perceiving, Sensing and Knowing* Los Angeles and Berkeley: University of California Press.
- Swinburne, R. (2001) *Epistemic Justification*, Oxford: Clarendon Press.
- Thau, M. (2002) *Consciousness and Cognition* Oxford: Oxford University Press.
- Travis, C. (2004) 'The Silence of the Senses' *Mind* 113, 57-94.
- Tye, M. (1984) "The Adverbial Theory of Visual Experience" *Philosophical Review* 93: (195–225).
- (1992) "Visual Qualia and Visual Content" in Crane (ed.) (1992).
- (1995) *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- (2000) *Consciousness, Color and Content* Cambridge, MA: MIT Press.
- Unger, P. (1975) *Ignorance: A Case for Scepticism* Oxford: Oxford University Press.
- (1984) *Philosophical Relativity*, Minneapolis: University of Minnesota Press.
- (1986) "The Cone Model of Knowledge" *Philosophical Topics* 14, 125-178.
- Valberg, J.J. (1992) *The Puzzle of Experience*. Oxford: Clarendon Press.

- Van Cleve, J. (1979) "Foundationalism, Epistemic Principles, and the Cartesian Circle" *Philosophical Review* 88, 55–91.
- Venn, J. (1876) *The Logic of Chance*, 2nd ed., Macmillan and co; reprinted, New York, 1962.
- Viger, C. (2006) "Is the Aim of Perception to Provide Accurate Representations? A Case for the 'No' Side" in R. Stainton (ed.) *Contemporary Debates in Cognitive Science*, Malden, MA: Blackwell Publishing, 275-288.
- Vinci, T. (1986) 'Critical Notice' *Canadian Journal of Philosophy* 16, 559-574.
- (1998) *Cartesian Truth*, Oxford University Press.
- Vintiadis, E. (2006) "Why Certainty is Not a Mansion" *Journal of Philosophical Research* 31, 143-152.
- Vogel, J. (1990) 'Cartesian Skepticism and Inference to the Best Explanation' *Journal of Philosophy* 87, 658-666.
- (2005) 'The Refutation of Skepticism' in Steup, M. and E. Sosa (eds.) *Contemporary Debates in Epistemology*, Malden, MA: Blackwell Publishing, 72-84.
- (2007) "Subjunctivitis" *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 134, 73-88.
- von Mises R. (1957) *Probability, Statistics and Truth*, revised English edition, New York: Macmillan.
- Williams, B. (1978) *Descartes: The Project of Pure Enquiry*, New Jersey: Humanities Press.
- (1983) "Descartes's Use of Skepticism" in M. Burnyeat (ed.) *The Skeptical Tradition*, Berkeley: University of California Press.
- Williams, M. (1986) "Descartes and the Metaphysics of Doubt," in A. O. Rorty (ed.) *Essays on Descartes' Meditations*, Berkeley: University of California Press.
- (1996) *Unnatural Doubts* Princeton: Princeton University Press.
- Williamson, T. (2000) *Knowledge and its Limits*, Oxford: Oxford University Press.
- Wilson, M. D. (1978) *Descartes*, London: Routledge & Kegan Paul.

- Wittgenstein, L. (1953) *Philosophical Investigations*, Malden, MA: Blackwell.
- Wright, C. (1985) "Facts and Certainty" *Proceedings of the British Academy* 71, 429–72.
- (1991) "Scepticism and Dreaming: Imploding the Demon" *Mind* 100, 87-116.
- (2000) 'Cogency and Question-Begging: Reflections on McKinsey's Paradox and Putnam's Proof' *Philosophical Issues* 10, 140-163.
- (2002) '(Anti-)Sceptics Simple and Subtle: Moore and McDowell' *Philosophy and Phenomenological Research* 65, 330-348.
- (2003) 'Reflections on the Acquisition of Warrant by Inference' in Nuccetelli, S. (ed.) *New Essays on Semantic Externalism and Self-Knowledge*, MIT Press, 57-77.
- (2004) 'Warrant for Nothing (and Foundations for Free)?' *Aristotelian Society Supplement*.