



University of Kentucky
UKnowledge

University of Kentucky Doctoral Dissertations

Graduate School

2011

CARTESIAN SKEPTICISM AS MORAL DILEMMA

Jennifer Woodward

University of Kentucky, kuoppala@mscd.edu

[Right click to open a feedback form in a new tab to let us know how this document benefits you.](#)

Recommended Citation

Woodward, Jennifer, "CARTESIAN SKEPTICISM AS MORAL DILEMMA" (2011). *University of Kentucky Doctoral Dissertations*. 809.

https://uknowledge.uky.edu/gradschool_diss/809

This Dissertation is brought to you for free and open access by the Graduate School at UKnowledge. It has been accepted for inclusion in University of Kentucky Doctoral Dissertations by an authorized administrator of UKnowledge. For more information, please contact UKnowledge@lsv.uky.edu.

ABSTRACT OF DISSERTATION

Jennifer Woodward

The Graduate School
University of Kentucky

2011

CARTESIAN SKEPTICISM AS MORAL DILEMMA

ABSTRACT OF DISSERTATION

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the College of Arts and Sciences at the University of Kentucky

By
Jennifer Woodward

Lexington, Kentucky

Director: Dr. Brandon Look, Professor of Philosophy

Lexington, Kentucky

2011

Copyright © Jennifer Woodward 2011

ABSTRACT OF DISSERTATION

CARTESIAN SKEPTICISM AS MORAL DILEMMA

I argue that despite the fact that there can be no strong refutation of skepticism it remains that ignoring skeptical hypotheses and relying on one's sensory experience are both sound epistemic practices. This argument comes in the form of arguing that we are justified in ignoring skeptical hypotheses on the grounds that (1) they are merely logically possible, and (2) the merely logically possible is rarely relevant in the context of everyday life. I suggest that (2) is true on the grounds that the context of everyday life is one in which our epistemic pursuit of truth is mixed with other pragmatic goals. The result of this mix is that the pursuit of truth can conflict with our goal of avoiding error in such a way that we must choose to prioritize one goal over the other.

The above choice implies that skepticism comes at an epistemic cost not acknowledge in the contemporary literature on external world skepticism. This epistemic cost of skepticism means that the relative risk of error involved in relying on sensory experience is not as epistemically problematic as has often been assumed. These considerations allow an anti-skeptical position in which relying on sensory experience is prima-facie justified despite the possibility of being a brain in a vat. In this paper I explore what such a position might look like and what the implications of such a view might be for relevant alternatives positions, the closure debate, and the concept of differing epistemic perspectives in contemporary epistemology.

KEYWORDS: Cartesian Skepticism, External World Skepticism, Relevant Alternatives and Skepticism, Epistemology and Skepticism, Skepticism

CARTESIAN SKEPTICISM AS MORAL DILEMMA

By

Jennifer Woodward

Dr. Brandon Look

Dissertation Director

Dr. Arnold Farr

Director of Graduate Studies

Monday May 2, 2011

Date

RULES FOR THE USE OF DISSERTATIONS

Unpublished dissertations submitted for the Doctor's degree and deposited in the University of Kentucky Library are as a rule open for inspection, but are to be used only with due regard to the rights of the authors.

Bibliographical references may be noted, but quotations or summaries of parts may be published only with the permission of the author, and with the usual scholarly acknowledgments.

Extensive copying or publication of the dissertation in whole or in part also requires the consent of the Dean of the Graduate School of the University of Kentucky.

A library that borrows this dissertation for use by its patrons is expected to secure the signature of each user.

Name

Date

DISSERTATION

Jennifer Woodward

The Graduate School
University of Kentucky
2011

CARTESIAN SKEPTICISM AS MORAL DILEMMA

DISSERTATION

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy in the
College of Arts and Sciences
at the University of Kentucky

By

Jennifer Woodward

Lexington, Kentucky

Director: Dr. Brandon Look, Professor of Philosophy

Lexington, Kentucky

2011

Copyright © Jennifer Woodward 2011

TABLE OF CONTENTS

Abstract.....	iii
Chapter One: Introduction.....	1
1.1: Introduction to External World Skepticism.....	1
1.2: The importance of knowledge and responding to the skeptic.....	9
1.3: Skepticism as a distinctly epistemic issue.....	20
1.4: Another look at the skeptic’s argument.....	27
1.5: Outline of remaining chapters.....	29
Chapter Two: Skepticism and Epistemic Perspectives.....	36
2.1: Laurence Bonjour and Epistemic Perspectives.....	36
2.2 Epistemic Perspectives and Skepticism.....	43
2.3: Nagel’s View From Nowhere.....	48
2.4: Audi and the <i>Structure of Justification</i>	52
2.5: Duncan Pritchard’s <i>Epistemic Luck</i>	55
2.6: Summary.....	60
2.7: Strategy for Handling the Skeptic.....	62
Chapter Three: Belief Formation in the Context of Everyday Life.....	73
Chapter 3.1: Introduction.....	73
3.2: Are there Relevant Disanalogies between Scenario A and Scenario S?.....	82
3.2.1: “Skeptical Alternatives are irrelevant because ...life”.....	86
3.2.2: Assumptions of reliability are <i>epistemically</i> justified.....	95
3.2.3: Agents are not blameworthy for relying on sensory experience...103	
3.3: Other Relevant Alternatives Views.....	104
3.4: “Assumptions of Reliability... are Truth-Conducive”.....	105
3.3: Conclusion.....	110

Chapter Four: Justifying Knowledge Ascriptions in Everyday Life.....	116
4.1: Characterizing the Problem.....	116
4.2: Henderson and Horgan’s Pascalian Wager.....	131
4.2.2: Grounding Assumptions of Reliability as Rational.....	138
4.3: The Pyrrhonian Objection.....	145
4.4: A Thought Experiment: The SK principle versus the AR principle.....	152
4.5: Conclusion.....	156
Chapter Five: Skepticism and Shifting Epistemic Perspectives.....	161
Chapter 5.1: Introduction.....	161
5.2: The Bigger Picture: How might a Theory of Knowledge Look?.....	162
5.3: Relevant Alternatives Conditions and “Ruling Out”.....	173
5.4: Lewis’s Rule of Actuality and BIV’s.....	178
5.5: Epistemic Perspectives and Closure.....	181
5.6: Other instances of mistakes in shifting perspectives.....	187
5.7: Conclusion.....	193
Chapter Six: Remaining Questions and Concluding Remarks.....	196
Chapter 6.1: Summary.....	196
6.2: Skepticism in a Philosophical Context.....	199
6.3: Moore and Skepticism.....	202
6.4: Sensory Experience, Knowledge, and Luck.....	205
6.5: Contextualism, Relevant Alternatives, and the Rule of Exclusion.....	206
6.6: Other Solutions to Skepticism.....	210
6.7: Concluding Remarks.....	212
Bibliography.....	219
Vita.....	222

“Introduction”

1.1: Introduction to external world skepticism:

Generally, external world skeptics believe that we do not possess knowledge of the contingent aspects of the world around us. Traditionally, such skepticism involves denying that sensory experience can ground knowledge. However, it is possible on this particular kind of skepticism to allow knowledge of the external world via means other than sensory experience. It might be possible, for example, to grant a priori knowledge of necessary truths via reason, and such a belief might constitute a kind of external world knowledge. If we look to the particulars of the argument that Descartes gave in Meditation I, concern about his external world beliefs stemmed from concerns about the reliability of the *source* of those beliefs; specifically, sensory experience. Descartes did not question the notion of external world knowledge *in principle*. Instead, he found that beliefs grounded in sensory experience fell short of counting as knowledge because the method by which the beliefs were formed admitted of too much doubt to yield knowledge.

Descartes notes early in his *Meditations* that:

All that I have, up to this moment, accepted as possessed of the highest truth and certainty, I received either from or through the senses. I observed, however, that these sometimes misled us; and it is the part of prudence not to place absolute confidence in that by which we have even once been deceived.¹

¹ See Descartes, *Meditations*, pg. 47.

In this moment Descartes set out that skeptical worry that has come to be known in the contemporary literature as Cartesian skepticism. While Descartes' skepticism is often thought of as grounded in the existence of uneliminated skeptical hypotheses, we can see in this quote that the fundamental question for Descartes was how we can be justified in relying on a source that has been found, at times, to be unreliable. Of course, Descartes felt that he could not justify such reliance; that an even occasionally unreliable source cannot yield real knowledge. This worry is particularly salient if we have no way for determining when the source in question is functioning reliably and when it is not. The possibilities of dreaming or an evil demon embody instances in which our source *seems* to be functioning reliably, when, unbeknownst to us, it is not. Descartes general strategy for analyzing the whole of his beliefs lies in evaluating sources of beliefs. In light of the possible failure of the senses, in light of the evil demon and dreaming scenarios, and the occasional de facto failure of the senses, new foundations must be found if we are to possess knowledge. The possibly and occasionally unreliable is not adequate to serve as a foundation for knowledge.

While many in the contemporary literature have focused on the above skeptical hypotheses, it is helpful to remember that at its heart Cartesian skepticism expresses the doubt about the ability of sometimes or *possibly* unreliable mechanisms to produce knowledge. If we reason along such skeptical lines we too might wonder how a source that might be massively unreliable can possibly yield knowledge. Even worse, we might begin to wonder whether or not it's rational or reasonable to rely on a source that we recognize might be deceiving us. The skeptic feels that since our sources have occasionally been locally unreliable, and might be globally and massively deceiving us,

we cannot possibly claim knowledge on such weak grounds. These are the questions particular to external world skepticism, and these are some of the questions one who wishes to reject skepticism must grapple with.

The concern over sensory experience in Cartesian skepticism seems to stem, primarily, from the existence of an array of skeptical alternatives such as the brain-in-the-vat or the evil demon scenarios. These skeptical scenarios expand the reason to doubt sensory experience from the classical examples of sensory illusions in which something looks one way but is another, to a worry about massive and general deception that can be neither confirmed nor refuted. Skeptical alternatives are particularly worrisome if we are concerned with avoiding error because they embody scenarios in which the error is general, massive, and cannot be detected through the methods available to us for detecting error.

This Cartesian skepticism undermines knowledge by attacking the justification that sensory experience provides for our beliefs. This kind of skeptic will argue that sensory experience can be mistaken, and that we cannot rule out the hypotheses that represent the possibility of error. The argument is that since we cannot rule out hypotheses in which our sensory experience is mistaken, we fail to have grounds for preferring our ordinary beliefs about the world over rival skeptical hypotheses. Given the above, the conclusion is that we fail to know what we take ourselves to know about the world. That is, we will fail to know such ordinary things as “I have hands”, and that there are trees in the world, and that there are other people in the world. This conclusion is so counter-intuitive that it tempts many to reject skepticism outright. However, as

tempting as this response might be, it is important that we have more substantive grounds than a mere dislike of the conclusion for rejecting the skeptic's argument.

The primary questions that come from this simplified skeptical worry are questions like how are we justified in relying on sensory experience when we have no *general evidence* that it is in fact reliable, how can we claim *to know* a hypothesis when we cannot rule out its competitors? Are our grounds for our beliefs the arbitrary and psychologically convenient choice the skeptic claims them to be? It is this particular, somewhat abstracted, skeptical puzzle that I have chosen, like many contemporary epistemologists, to focus my efforts on. And I, like many epistemologists, am concerned with how the contemporary internalism and externalism debate is affected by this particular skeptical issue in epistemology.

The current debate on internalism and externalism seems best characterized as a debate between epistemologists over whether or not that which serves as justification for knowledge must be accessible to the agent via introspection or not. While internalists contend that justification must be accessible to the agent's perspective, externalists deny that the accessibility requirement is necessary for knowledge. Externalists are willing to count beliefs as knowledge even when that which justifies the agent's belief is not accessible to the agent herself. For many externalists, as long as the agent's belief was *in fact* formed in a reliable way, the agent counts as knowing whether or not the agent is *aware* of the reliability of her belief forming mechanism. On other externalist views, it might be that an agent exhibits certain epistemic virtues the possession of which confer the status of knowledge on her beliefs, regardless of whether or not she is aware of her virtuous epistemic practices *qua* virtuous epistemic practices. Some externalists count

the latter as knowing that one knows, and make sharp distinctions between knowing that you know and mere knowledge, between knowledge and knowledge ascriptions. Thus, we might think of the internalism and externalism debate as one in which internalists require the accessibility of the justifiers and externalists do not.

While I will not address internalism and externalism in any explicit way in the argument to come, there are implications for this debate that can be extrapolated from the argument given. I must admit early on that the implication of approaching this issue from such a narrow standpoint is that my discussion will lack much of the detail and historical finesse that a comprehensive discussion of skepticism would involve.

My argument takes Duncan Pritchard's work in *Epistemic Luck* as a starting point and focuses on the issue of *knowledge ascriptions* and skepticism. Pritchard comes to the conclusion in this work that externalism goes some way towards addressing skepticism for knowledge but that such a solution leaves skepticism intact with regard to our knowledge claims. My work here starts with this latter concern about skepticism as it applies to knowledge claims.

I utilize Bonjour's concept of Epistemic Perspectives to both justify a strong distinction between the conditions for knowledge and knowledge ascriptions as well as provide a justification for treating skeptical hypotheses as irrelevant in most everyday contexts. Because I focus primarily on knowledge ascriptions, and there tends to be more agreement on this score between internalists and externalists, I have aimed my discussion in a way that is meant to appeal to both sides of this debate. While it may be that knowledge need not require access to the justifiers, as externalists contend, it seems less plausible to deny the accessibility requirement of internalism for knowledge ascriptions.

That is, it seems fairly counter-intuitive to allow one to truly *claim* knowledge without having any access to, or awareness of, the justifiers. Duncan Pritchard is an example of an epistemologist who is an externalist with regard to knowledge but seems to require accessibility for knowledge ascriptions.

I will argue that skeptical hypotheses are generally only relevant in a philosophical context and that we have good epistemic reasons for treating them as irrelevant in most everyday contexts. The aim of this argument is to defend the practices of everyday knowers without appeal to fancy philosophical argument. Thus, my argument draws out and relies upon common sense attitudes that reject skepticism and defends those attitudes as epistemically justified. The appeal of this approach is that it results in the defense of assumptions we make about sensory experience in a way that justifies everyday knowers without appeal to a line of reasoning that it is unlikely everyday knowers have ever appealed to. The line of reasoning given in the course of the argument to come could be applied more generally to justify assumptions of reliability for other methods of belief formation such as testimony and memory.

In “Skepticism and Rationality” Richard Foley argues that it is a brute fact about epistemic endeavors that they are fallible, and that we too, as possible knowers, are fallible, and that our methods of belief formation themselves are inherently fallible. He states:

To be sure, this involves a leap of intellectual faith. It involves our having confidence in those intellectual methods that are deeply satisfying to us despite the fact that we cannot vindicate this confidence in a non-question begging way. This may be regrettable but it is also undeniable. The reality of our intellectual lives is that we are working without nets. No procedure, no amount of reflection, no amount of evidence gathering can guarantee that we won't fall into error, perhaps even massive error. We are thus forced to choose between moving forward in a way that we, upon

reflection, would take to be effective and proceeding in a way that we would not take to be effective. If we are rational, we opt for the former.²

In his *Epistemic Luck* Duncan Pritchard talks about an “Epistemic Angst” that remains for our internalist justifications and for our knowledge claims even in light of a successful externalist solution to skepticism with regard to *knowledge* conditions.³ This epistemic angst is characterized by a necessary leap of faith; we are akin to “acrobats” who are not certain that the safety net that appears below is really there.

These views articulate a certain acceptance of our inherent fallibility, and they articulate what our epistemic endeavors look like amidst such fallibility. It is difficult to deny such fallibility, and in the argument to come I will embrace this portrait of our epistemic lives. We certainly must embrace a sense of fallibility for knowledge if we wish to maintain that such paradigms as science do in fact achieve and attain knowledge. Given the technologies and advances we have made on the heels of these epistemic endeavors it seems hard to account for such success if the fruit of that labor does not count as knowledge. Yet, the nagging skeptical worry remains. Indeed, the skeptical worry continues to have some bite to it.

The skeptic might be right to wonder if we are justified in claiming something as strong as *knowledge* in light of such profound fallibility. Why is taking a *leap of faith* the right epistemic response to such fallibility? Why shouldn't the acceptance of our fallibility humble us in such a way that we refrain from claiming knowledge? The skeptic might be right to wonder why such views act as though knowledge claims are justified even though they acknowledge that we might be in massive error. Such a

² See Foley, pg. 331.

³ Duncan Pritchard, *Epistemic Luck* (Oxford: Oxford University Press, 2005), 225-253.

skeptical worry questions whether or not it is *rational* to claim knowledge in light of uneliminated possibilities of error. It is these skeptical worries that hold some weight and deserve further answer. It is not so clear that the brute fact of fallibility implies that knowledge claims continue to be rational after we have recognized the myriad ways in which we might be wrong about some of our most basic external world beliefs. Our stance towards sensory experience needs further justifying, and it is my sense that such justification can be given. Thus, in the argument to come I will be looking to articulate how and why we continue to be justified in claiming knowledge in light of our inherent fallibility.

It is important to note what I intend to accomplish in the coming arguments. I intend to show that endorsement of sensory experience is not an instance of epistemic failure or an instance of irrationality. For the skeptic, if we do not have sufficient reason to believe, then we are irrational if we do so. I intend to show that this is not always the case; and it is not the case with regard to the general reliability of the senses. What I am aim to show is not that non-skeptics are epistemically right and skeptics epistemically wrong. Rather, I merely wish to deny that *non-skeptics* are somehow *irrational* for believing on what count as less than sufficient grounds for the skeptic. Thus, what will come out of the arguments to come is that there is an *epistemic* stale-mate between the skeptic and non-skeptic. Thus, the distinction between these views lies in the epistemic *values* each embraces and embodies. These values are foundational in the sense that they are not epistemically driven by the evidence.

The position of skeptic or non-skeptic reflects distinctions in where each places value for epistemic “goods and bads”. I will argue that the best way to think of

skepticism is as an epistemic approach that endorses very high standards for knowledge because of a concern to avoid error. The non-skeptic, on the other hand, or the mitigated skeptic, places value on truth while *also* placing value in avoiding error. The evaluation of the skeptical versus non-skeptical positions needs to occur as an evaluation of the epistemic values they embrace. In this light, I believe it can be shown that the non-skeptic's position is at least as *epistemically appropriate* as the skeptic's. Thus, I am aiming not to refute skepticism, but to take the sting out of some of the worries about rationality that arise out of taking skepticism seriously.

1.2: The importance of knowledge and responding to the skeptic:

We might ask ourselves, what is the importance of knowledge and why is it important that we defend it against skepticism? Our first endeavor is to justify spending time on responding to skepticism in the first place. Indeed, it may be tempting given the complexity of the problem to dismiss skepticism as a mere parlor trick rather than as a serious argument with an unlikable conclusion. The strategy of dismissing skepticism as a sort of parlor trick is unappealing given the lack of explanation for its status as a mere trick, rather than as serious philosophical argument. Surely, we do not want to endorse methodologies that allow us to throw off the requirements of careful argumentation whenever we see fit. Giving a serious explanation about why skepticism is not a threat to knowledge amounts to a serious treatment of and response to skepticism. Nonetheless, it may be helpful to explore in more detail what is the importance of knowledge as opposed to justified belief or true belief and what is the importance of responding to skepticism.

Since knowledge of the world is the focus of external world skepticism, I will focus my discussion on the value of external world knowledge. Further, I will take it that I am being asked to discuss the value of knowledge as opposed to merely justified belief, or as opposed to merely true belief. Roughly, to possess knowledge of the world is to be in an epistemically ideal state with regard to our cognitive relationship to the world. When we possess knowledge of the world we do so reliably and consistently rather than accidentally or haphazardly. Whether or not we regularly *know when* we possess knowledge, possessing knowledge implies that we have a better cognitive relationship to the world than possessing merely justified belief.

In answering concerns about the value of knowledge, I do not want to restrict myself to the first person perspective. The first person might be characterized as the perspective in which we are generally making and assessing knowledge claims. Admittedly, from this perspective we may often *function* practically on the basis of the *justifications* for our beliefs. So for example, we *claim* to know often on the basis of the evidence we have for our belief, we feel entitled to communicate our beliefs on the basis of the evidence we have, etc. These practices make it look as though justified beliefs are all we need to navigate through epistemic practices in everyday life. In the course of this discussion I aim to make two general points: (1) attempt an answer at why knowledge is preferable to merely justified belief and why knowledge is preferable to merely true beliefs, and (2) establish that if the skeptic is right, not only do we fail to possess knowledge, we fail to possess any degree of epistemic justification for our beliefs. Establishing (2) is important to blocking the common argument that the skeptic employs unusually high standards of knowledge, and so it is possible to concede that we fail to

know in the skeptic's sense but still possess some degree of epistemic justification for our beliefs.⁴ This move fails unless the *requirements* on what counts as evidence are reconceived. Unless we address the skeptic's views on evidence we have to concede both that we fail to know, and that we fail to possess *any* degree of epistemic justification for our beliefs. Further, many have found the skeptic's standards and her argument as having some fundamental appeal. Thus, we will need an explanation and understanding of how and why her standards are too high.

Regarding the value of knowledge as opposed to justified beliefs, it is worth pointing out that justifications do not logically entail the truth of our beliefs, the connection between justification and truth is not that close.⁵ It seems that the connection between our justification and beliefs is to a large degree, at least when it comes to knowledge about the world, contingent. To this extent, our justifications may indicate the truth of our beliefs given that certain contingent facts about the relationship between the world and our cognitive faculties hold. However, our beliefs about the world will often amount to knowledge when a truth producing relationship holds between the world and our cognitive faculties. Thus, the value of epistemic practices involving justified beliefs will depend for their rationality on the same contingent factors that knowledge depends upon. For these reasons, it does not behoove us in the skeptical debate to transition from

⁴ This amounts to the common line that the skeptic operates with unusually high standards and the only response we need to give to the skeptic is to simply say that we are interested in a level of justification that is lower than what the skeptic is operating with.

⁵ In *Putting Skeptics in their Place*, John Greco suggests that the relationship between our epistemic justifications and our beliefs is not even quasi-logical, our justifications neither deductively nor inductively imply the truth of our beliefs, and that this is one of the lessons of skepticism. He argues, further, that if we are to address skepticism we must rethink the relationship between evidence and beliefs.

a debate on knowledge to justified belief. As it stands, it seems to me that the skeptic attacks knowledge by attacking justification and that it is the justification element that is hard to defend against skepticism. For this reason alone, it will not work to grant the skeptic knowledge while trying to maintain that we have epistemically justified beliefs even if they fail to amount to knowledge. The skeptic aims to undermine both knowledge and justification.

Knowledge is a better epistemic state to be in than the possession of an epistemically justified belief simply because the possession of knowledge implies the belief in question is true, whereas the possession of even an epistemic justification may not.⁶ There are two senses of justification I'll be using and what I am calling *epistemic* justifications are justifications that do in fact increase the likelihood of the truth of one's belief.⁷ By *justified* belief I mean a belief that from the agent's perspective is likely to be true given her other beliefs and her evidence. The latter sense of justification amounts to something like whether or not it is understandable, or rational from the agent's own perspective, to believe a certain proposition. As I have spelled out epistemic justification I cannot see, in the face of skepticism, any reason to prefer talking about it as opposed to knowledge as some kind of concession to the skeptic.

⁶ This may well depend on how one spells out epistemic justification, and whether or not requirements about defeaters are placed on one's evidence in order for it to count as justification. These are not issues I want to go into, so I am talking about epistemic justification in the sense that it is either a reason or evidence that indicates the truth of one's belief without any consideration of defeaters. This strikes me as okay, because defeaters are often spoken of as a fourth condition on knowledge as opposed to being built in at the level of justification. This is good, as defeaters may not be accessible from the first person perspective, and justifications are supposed to be.

⁷ I have spelled out epistemic justification in a particularly externalist way, such that whether or not an agent possesses an epistemic justification may not be something the agent can access from the first person.

As far as the latter sense of justification goes it seems clearly less valuable than knowledge. Whereas both epistemic justification and knowledge depend upon the skeptic's being wrong, the possession of merely justified beliefs does not. We may be justified in our beliefs even if we are unable to access the external world, and none of our justifications actually imply any likelihood that our beliefs are true. For these reasons, it is much more preferable to possess knowledge or epistemic justification than it is possess merely justified beliefs. Epistemic justification counts towards truth in some way, and knowledge implies it; since we are aiming for truth in the first place we are aiming to possess one of these two, and knowledge is clearly preferable because of its' connection to truth.

A quick example may help make the point. If an agent looks down, has the sensory experience of seeing her hands, and comes, on this basis, to believe she has hands; she is justified in this belief whether or not she is a brain in a vat. However, if we imagine that this particular agent is a BIV there is a sense in which even though she is justified in her belief her epistemic state is deprived- she possesses neither knowledge nor epistemic justification. Further, her belief that her sensory experience indicates the *truth* of her belief is false. Her belief that she possesses an epistemic justification for her belief is false, and, of course, her belief is itself false. Her justification doesn't indicate what she takes it to indicate – that she has hands. On the other hand, the possession of epistemic justifications and knowledge do imply that we are successful in accessing the external world. In so far as this is important to us, and I think it is, knowledge is important to us. Knowledge implies the successful attainment of our most prized epistemic goal- namely, the attainment of truth.

Even more than the above, it seems that epistemic justifications and knowledge are what we are really interested in. When we communicate our beliefs to one another we are concerned with whether or not the belief is true, not with whether or not it is understandable that the agent holds the belief. It may well be that if we genuinely believed knowledge of the world impossible, then we would not treat justifications as epistemically relevant. We use justifications to show why we think our particular belief amounts to one of those produced by a good cognitive relationship to the world. So in some ways, these other practices are parasitic on our having a cognitive relationship to the world that makes knowledge possible, and does in fact often produce knowledge. While brains in vats may well justify their beliefs to each other in much the way we do, the depravity of their epistemic situation makes it such that their justifications are not epistemic justifications, and these are ultimately the ones we care for the most.

To summarize, knowledge is valuable as the primary product of agents having a good cognitive relationship to the world. All of our epistemic practices depend, for their rationality, on this relationship holding. If such a relationship fails to hold, then it ultimately does not matter what we believe about the world and why. Agents in skeptical worlds are essentially cut off from their real environments. If the argument that knowledge is unimportant is that whether or not we possess it depends on things we cannot access, and we seem to be able to function only with those things we *can* access from a first person perspective, then these concerns hold not only for knowledge but for any sense of justification on which a justification is something that does in fact connect one's belief with the truth.

It has also been argued that knowledge should not be the focus of the epistemologist's inquiries because it is not a single, unified, analyzable concept. All I want to say in response to this is that we need not claim that knowledge is a single unified concept in order to maintain that the concept/s are important. Further, if we conclude that analyses of knowledge are unsuccessful, we can still utilize the concept in our epistemological projects.⁸ Further, even if it is the case that the concept knowledge is unanalyzable, this would not be a reason to favor *justification* as it seems as difficult to spell out the conditions for epistemic justification as for knowledge (indeed, it is *this* aspect of knowledge that is difficult to analyze).

Nonetheless, one may still make a case that epistemologists could focus only on justification rather than knowledge. The argument might go roughly as follows: from a first person perspective we just do not know (and there's no way to find out) whether or not we have a good cognitive relationship with the world. Thus, the best we can do is to *assume* that we do and produce justifications that would indicate the truth of our beliefs *if* such a relationship held. I think this argument is perfectly fine when our focus is a first person perspective, and justifying our beliefs from this perspective. It may well be that there is no way to confirm that the world and our minds are such that knowledge is possible, or that our justifications are what we take them to be – epistemic justifications. However, I have two points to make in response to this argument: (1) this seems no more devastating for knowledge than it is the practice of giving and accepting justifications for our beliefs, because the rationality of these practices depends as much as knowledge does

⁸ Timothy Williamson does this in his book *Knowledge and Its Limits*. He argues that knowledge is a basic concept, and is more fundamental than justification. See Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000).

on whether or not we have reliable access to the external world, and (2) even if it were true that focusing on justification for epistemic questions in the first person makes sense, this does not undermine that from a third person perspective, knowledge is the better epistemic state because in order to be in this state we must actually be in a good cognitive relationship to the world.

I will contrast the benefits of possessing knowledge with that of merely possessing true beliefs only briefly before discussing the importance of the success or failure of skeptical arguments. The possession of knowledge is more valuable than the possession of merely true belief, because the possession of knowledge implies a successful connection between our cognitive faculties and the world that the possession of true belief does not. The possession of true beliefs need not inspire confidence in our ability to attain further true beliefs, nor inspire confidence in our abilities to successfully grasp reality *generally*. Granted, our ability to function with true beliefs may be similar to that of our ability to function with knowledge, but only at a very basic level. If I wanted to go to the grocery store and I accidentally, but truly came to believe there is a grocery store on 2nd street, I may well be as successful locating and utilizing the grocery store as I would be if I had known there was a grocery store on 2nd street. However, Timothy Williamson makes a legitimate point in his *Knowledge and Its Limits* when he argues that our beliefs that amount to knowledge are safe.⁹ They are beliefs that we are more likely to hold onto in the face of evidence to the contrary. In a similar sense, beliefs formed in a way that is not generally successful but only accidentally truth yielding in a case or two are not likely to be beliefs that inspire the kind of confidence necessary to avoid being

⁹ Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), 62.

easily misled by evidence to the contrary. For these reasons, the possession of knowledge is better than the possession of true beliefs.

It may be thought that we need not care as to whether the skeptic is right because the skeptic's demands are so unreasonable that they are impossible to meet, so who cares if we don't? At this point it may be worth reviewing the structure of skeptical arguments. Where H is some skeptical hypothesis and O is a proposition one would normally take herself to know, the following is a characterization of the structure of skeptical arguments borrowed from Keith DeRose¹⁰:

1. I don't know that not-H.
 2. If I don't know that not-H, then I don't know that O.
- C. I don't know that O.

The first premise is established by our usual inability to rule out skeptical hypotheses based on the evidence we possess. If we recall skeptical hypotheses such as the hypothesis that one is currently dreaming or the hypothesis that one is a brain in a vat, they are hypotheses that are consistent with *any* sensory experience we can produce in favor of our ordinary beliefs about the world.

For example, given the above structure if H is the skeptical hypothesis that S is a brain in a vat, and S claims to know that she has hands (O), her evidence for O is her sensory experience of seeing her hands. However, her sensory experience of seeing her hands does not rule out the possibility that she is a brain in a vat, since brains in vats also have sensory experiences of seeing their hands. Further, it is true that if she is a brain in a

¹⁰ Keith DeRose, "Introduction: Responding to Skepticism," in *Skepticism: A Contemporary Reader*, ed. by Keith DeRose et al (Oxford: Oxford University Press, 1999), 2.

vat, then she fails to have hands, thus she fails to know that she has hands. So the reasoning goes that since S's evidence for O is consistent with skeptical hypothesis H, S is unable to rule out H (premise 1), and since she is unable to rule out H, she does not know O.

On the skeptic's standards in order to count as knowing O one's evidence must rule out all hypotheses incompatible with O. Yet, there are many hypotheses consistent with the evidence one has for O and yet incompatible with it as well, thus for any ordinary belief about the world we fail to count as knowing it. Another way to put it is to say that the skeptic operates with the following epistemic principle: if the grounds for one's belief that p is evidence e, then one is not entitled to believe anything stronger than what e indicates. If e is consistent with a set of possibilities, then one is entitled to the belief that *one* of the set of possibilities holds (given the assumption that we inhabit *a* particular world), but not entitled to believe (at least not on the basis of e) that any *one* of the particular possibilities holds. In other words, one is entitled to believe that one of the possibilities is the actual one, but not entitled to infer that any given possibility is the actual one.

The skeptic's standards are so stringent that there may be the inclination to say, if *that* is what it takes to possess knowledge, then who cares if we don't have it. The skeptic demands our evidence rule out every logical possibility that is incompatible with our belief, yet this seems overly strict. Indeed it seems that we are rarely in such an ideal state, and yet we manage just fine. Further, we may even have the intuition that our inability to rule out ridiculous logical hypotheses fails to establish an epistemic

deficiency in either our evidence or the rationality of our belief. While I am sympathetic with this line, in a way I will articulate below, I think this kind of response, as it stands, is misguided.

Implicit in the above reasoning is that the skeptic operates with ridiculously stringent standards on knowledge, and it doesn't matter whether or not we reach those high standards because we often reach slightly lower standards in the justification of belief, or knowledge, and we manage just fine operating at the levels we do obtain. However, this line seems unsuccessful because the skeptic attacks our knowledge by attacking our evidence for our beliefs. As such, the skeptic undermines not only our knowledge but the justification for our belief, and perhaps even (assuming a voluntarist account of belief) the *rationality* of our beliefs.¹¹ If part of belief is the stance or attitude that the belief in question is true, then, if the skeptic is right, this stance is entirely unjustified. Further, it would seem to be a cognitive defect on our part that we fail to appreciate that our evidence fails to establish or support the particular beliefs we hold. Indeed, if the skeptic is right the only beliefs we might be entitled to on the basis of our sensory experience are disjunctive ones (either I have hands or I am in a skeptical world in which case I fail to have hands, etc).

If the skeptic is right, our epistemic situation is much worse than our lacking knowledge, the skeptic has ultimately questioned the *rationality* of our beliefs. If we concede the skeptic's point, and continue to maintain our ordinary beliefs, we must

¹¹ In "Contextualism and the Problem of the External World," Ram Neta argues that the skeptic raises the standards for evidence not the standards for knowledge.

follow Hume in adopting psychological and pragmatic justifications for our beliefs, but we must admit that there are no *epistemic* grounds for our beliefs. So it is not *just* as though we fail to possess knowledge if the skeptic is right. Rather, if the skeptic is right, we fail to possess *any* degree of justification for thinking that our beliefs are true.¹² So, to conclude, it is not just knowledge that is at stake with the skeptic; rather, the rationality of all our epistemic practices is undermined if the skeptic is right.¹³

1.3: Skepticism as a distinctly epistemic issue:

In the previous section I articulated the reasoning behind taking skepticism seriously. Additionally it should be noted that I will be approaching skepticism specifically as an *epistemic* issue. I will work within the confines of the epistemic aspect of skepticism. There are a few recent responses to skepticism that are in some sense distinctly non-epistemic. That is, they address the skeptical issue by addressing some necessary but non-epistemic assumptions on the part of the skeptic. I think of these responses as rejecting the framework for skepticism that the skeptic has set up. While I do not deny the possibility of such responses, I have minimal interest in addressing the efficacy of these responses. The reason is that to my mind even if such a response pre-emptly the skeptical problem, thereby providing a solution, epistemic worries remain in

¹² This needs a little qualification as it is a bit too strong as is. Our beliefs would possess a higher likelihood of truth than beliefs that are inconsistent with our sensory experience, they would not however possess any higher likelihood of truth than any other hypotheses within the set of hypotheses consistent with our sensory experience, i.e. our beliefs would no more likely be true than skeptical hypotheses. It should also be pointed out that Bonjour gives an argument in which he attempts to show that our ordinary beliefs are more likely to be true than their skeptical counter-parts. This is an argument that we will look at later.

¹³ In light of these points, I am much more sympathetic to Williamson's position that justification cannot be separated from knowledge, and is the more basic concept.

the face of such responses. Therefore, part of taking skepticism seriously is to take seriously the idea that the skeptic has gotten a hold of something substantive and distinctly epistemic. This is so even if her conclusion is not fully warranted.

Despite the above, there is a need to say more about such responses. The two primary responses I have in mind here are a move to block skepticism via semantic externalism and an attempt to block skepticism via various positions on the nature of perception. Both of the above responses have evolved many distinct forms and have generated distinct sub-topics in skepticism. Since our primary concern is skepticism and its implications for the internalist and externalist positions in contemporary epistemology, I will be addressing the above responses in their most basic form.

Hilary Putnam initially and most persuasively argued that skepticism can be blocked on the basis of semantic externalism. Putnam advanced the argument that the statement “I am a brain in a vat” is a self-refuting statement.¹⁴ The reason is that the conditions that would make such a statement true would also change the meaning of the statement making the statement false. In other words, it is impossible for a brain in a vat to actually refer to things like brains and vats. Putnam states that “Although people in that possible world (referring to brains in vats) can think and ‘say’ any words *we* can think and say, they cannot *refer* to what we can refer to.” Thus there is the peculiar conclusion that there is no possible world in which the statement “I am a brain in a vat” will turn out true. We, therefore, have the reason to reject the brain in a vat hypothesis on the grounds that it is necessarily false.

¹⁴ Hilary Putnam, “Brains in a Vat,” from *Skepticism: A Contemporary Reader*, ed. Keith DeRose et al (Oxford: Oxford University Press, 1999), 31-32.

Any actual brain in a vat cannot refer to brains and vats and therefore cannot truly utter that statement. In the brain in a vat's world the words "brains" and "vats" will refer to whatever usually causes the internal sensory experience associated with brains and vats, but in a brain in the vat world it won't actually be brains and vats that causes utterances of "brains" and "vats". So this move rejects the problem of skepticism on purely semantic grounds.

What is interesting about this move, and the many articulations of it, is that it addresses skepticism without addressing what I take to be the core epistemic issues touched on by skepticism. In other words the above response does nothing to assure us that we are not actually brains in vats, or that we are not massively deceived, or that knowledge is possible in the face of the possibility of massive error. And while it may insure that our beliefs are (by and large) true, it does so in a superficial way. It is superficial in my mind, in the sense that our beliefs may be true, and, yet, if we were brains in vats we would have absolutely no knowledge of the nature of what our words refer to, or, for that matter, would we possess any knowledge of the nature of the world we live in. We would be uttering true statements but not have any genuine understanding of the world around us. Such a deficiency is surely an *epistemic* one and signals to my mind that the response from semantic externalism does not respond to this core epistemic issue. Thus, to my mind the skeptical problem remains even in light of such an interesting attempted solution.

The second response that side-steps skepticism is a response based in a rejection of the skeptic's view of perception and sensory experience. Many skeptics seem to be worried over what many have called the problem of the criterion, or the problems

associated with such a representationalist view of perception as that seen in Locke's writings. In the *Meditations* Descartes goes from worrying about fairly innocuous instances of error such as the occasional instance in which he views an object from a great distance and incorrectly infers various properties of an object, to more *widespread* concerns about error, such as the possibility that he is dreaming and the possibility of an evil demon. At root in these worries is the assumption that sensory experience is somehow detached from what it purportedly represents, thereby introducing a possibility of error. If sensory experience is merely representational, then we might wonder how we know that it faithfully *represents* in the way we take it to.

Of course, we recognize that our beliefs are not beliefs about sense-data; rather they extend beyond subjective and mental sense-data to the external world. Our beliefs are aimed at what we take to be mind independent objects rather than the ideas of such objects. Such a view of sensory experience and perception inevitably leads to questions about whether or not sense-data is representational in the way we think it is. Since there is no way to get outside ourselves and access the external world in a way other than sensory experience, we have no way to check and confirm that sensory experience proximally represents the world in which we live. The evil demon scenario is one among countless possibilities in which sensory experience is misleading. Since there is no way to confirm the reliability of sensory experience skeptics will argue that we have no justification for relying on it as a method of belief formation, and we, thereby, are not justified in beliefs that are founded in sensory experience. In other words, we will not be justified in the vast majority of our ordinary beliefs about the world because they involve

an inferential leap from sensory input to the external world that is not epistemically justified.

If we question this approach to perception and replace it with a direct theory of perception in which sensation is not merely representational in its nature but is, instead, some kind of direct contact with objects in which the content of belief is not inferred or indirect, but directly contained in the sensory experience, then we have closed the gap between sensation and belief. Without the gap, the skeptic does not have reason to worry about whether or not sensory experience does indeed represent the world. Nor does she have grounds to question the inferential leaps made from sensation to beliefs about the world, for there are no inferential leaps.

Bonjour uses something like the above as a response to skepticism. The view of sensation that he puts forward grounds his foundationalism with regard to sensory experience. In his *Epistemic Justification*, he states, “Thus contrary to many recent critics of foundationalism, the idea that reality is in some circumstances simply *given* to the mind in a way that makes the truth of claims about it directly and unproblematically apparent is, after all, not a myth!”¹⁵ When seen as a response to skepticism, it seems best to treat this kind of response as rejecting the first premise of the skeptic’s argument; the premise that states “I don’t know that not-H”, where H is some skeptical hypothesis. The second premise of the skeptic’s argument is that if we don’t know that not-H, then we don’t know that O. This premise is a distinctly epistemic one and seems to have no relationship to the issue of perception. However, the skeptic establishes the first premise on the basis of her interpretation of our lack of evidence against H, and her claim that our

¹⁵ Laurence Bonjour and Ernest Sosa, *Epistemic Justification* (Oxford: Blackwell Publishing, 2003), 75.

evidence equally favors O and H. This is precisely what the direct perception theory attempts to deny. Our evidence for O is our sensory experience, and it is somehow directly connected to the belief O. Direct perception will deny that there is an *inferential* leap between the sensory experience that grounds O and the belief that O. Further, the direct perception theory will maintain in some way that that the content of O is somehow directly apprehended in the sensory experience that grounds O. Thus, on such a view there are in fact reasons to rule out H. Such reasons will be grounded in sensory experience that distinctly favors the content of our ordinary belief O over its' skeptical rival H.

While I do not want to claim to have reason to dismiss such responses with such a short and generic discussion of this substantive issue, I do wish to present some reasons for side-stepping this issue in the discussion to come. Two things strike me about the above response via direct perception theory. First, there is something that seems correct about the response. The skeptic operates with and seems to get her argument going on assumptions about perception that may well be worth questioning. In fact, in everyday life we seem to operate with a view of perception akin to direct or naïve realism that the above positions attempt to do justice to. The second thing is that the sensory experience would have to be quite loaded in terms of cognitive content in order to fully get around skeptical worries about the veracity of sensory experience, and this seems problematic and unlikely. There would have to be a very tight connection between the content of our beliefs and what is directly apprehended in sensory experience. In fact, the connection would need to be so tight as to almost provide certainty in order to avoid skepticism this way. If there is any gap between what is apprehended in perception and the cognitive

content of our ordinary beliefs about the world, then the skeptic's argument can get going.¹⁶ She can again imagine that sensory experience is misleading and that we make inferential leaps from the experience to beliefs that are unwarranted.¹⁷ Lastly, it seems to

¹⁶ Admittedly, Bonjour notes this worry in his argument and tries to show that if the fundamental properties of an object are directly apprehended in experience, then that alone seems sufficient to ground a preference for our ordinary beliefs about the world. This move seems plausible but it involves an assumption that fallible and less-than-certain evidence is sufficient. However, his view does not articulate how or why we might claim to know in light of the possibility of error that persists even on his view. This latter move is what is doing the bulk of the anti-skepticism work on his view, and, yet, it is not articulated explicitly nor defended. To my mind, it is this latter position – the view that less than definitive evidence is fine enough that has to be defended on epistemic grounds if we are going to refute skepticism.

¹⁷ I have Hume's skeptical arguments in mind here in particular, because he particularly goes after those gaps between sense data and the content of our ordinary beliefs about the world. Also, I want to note that I do not mean to *assert* an inferential relationship between beliefs about the world and sensory experience. In fact, it seems unlikely that the relationship is inferential; rather the relationship between sensory experience and beliefs about the world seems to be causal. I am speaking this way though, because skeptics such as Hume have treated the relationship as inferential and questioned whether the inferences are justified. Further, it does not seem to help any to point out that the relationship seems to be causal because it is difficult to see how such a causal relationship could confer justification in a way that is of interest to the skeptic.

me that such a tight connection between what is apprehended or given in sensory experience and the content of our beliefs does not seem to fit with ordinary experience. We do seem, in many instances, to be wary of sensory experience. This indicates that there is not such a tight relationship as would be needed to circumvent skepticism via the above kinds of views on perception.

The core of the skeptical problem is the possibility of not knowing the nature of the external world; the possibility of being grossly in error about our environment. This possibility is in itself deeply worrisome to most people and poses an epistemic problem because it challenges the idea that knowledge is possible in the face of such fallibility. While I won't argue for this here, I will venture to say that in my mind a solution to skepticism needs to come in the form of explaining how knowledge is possible in the face of the possibility of such error and needs to explain why we are justified in assuming that we are not brains in vats. What I am interested in is the fundamental *epistemic* worry that skepticism raises and possible epistemic responses to the problem. Therefore, I will largely be working within the traditional epistemology framework despite the relatively recent appearance of solutions to skepticism that step outside that framework.

1.4: Another look at the skeptic's argument:

Skeptics establish their claims by attacking the evidence that our claims to know rest upon. It is this aspect of the skeptic's argument that makes skepticism relevant to the contemporary and ongoing internalism/externalism debate in epistemology. If we go

back to Keith DeRose's characterization of skeptical arguments where H is some skeptical hypothesis and O is an ordinary belief, then the skeptic argues that we do not know on the basis of our inability to rule out H. We cannot rule out H because our evidence is consistent with it. Thus, according to the skeptic we have a number of hypotheses consistent with the evidence and no reason for preferring our ordinary beliefs about the world over the skeptical hypothesis H.

We should note that in such a situation, often the appropriate move is to suspend belief and wait until the evidence gives us reason to prefer one hypothesis over another. However, we should also note that in many such circumstances we widen our body of evidence to include pragmatic reasons as well as other non-epistemic factors. One of the things I will argue is that the skeptic's claim that we do not have reasons for preferring our ordinary beliefs over their skeptical competitors is not as strong as it initially seems. Further, if we take an externalist approach to knowledge, then internalist justifications no longer play the role of the primary grounds for knowledge. Instead, they play a role as the grounds for knowledge *claims*. This slightly weaker epistemic role allows us to provide somewhat pragmatic justifications for preferring our ordinary hypotheses over their skeptical competitors that would not be strong enough to ground knowledge, but are strong enough to ground knowledge claims. This view will be outlined in chapter 5.

Admittedly, the skeptic is simply right in her argument that our evidence, as long as it is restricted to sensory experience (and characterized in a subjective fashion), is perfectly consistent with any number of skeptical hypotheses. If our challenge is to find the evidence that shows we are not brains in vats, then the skeptic has set forth an impossible task that cannot be fulfilled. Rather than look for *evidence*

that justifies us in preferring our ordinary beliefs about the world, we will look at various reasons we might have for preferring our ordinary beliefs about the world over skeptical hypotheses.

Arguments for skepticism that occur from deductive closure seem, to my mind, to involve the issue of evidence. Briefly summarized, skepticism based on deductive closure may be characterized as Stine did in her work on skepticism and closure, “The skeptical argument goes: If you know it is a zebra, and you know that it being a zebra entails it's not a painted mule, then you know it is not a mule painted to look like a zebra. But you do not know the last, so you do not know the first- i.e., you do not know it is a zebra.”¹⁸ While there are unique issues revolving around closure and skepticism, the skeptical argument based on closure still rests on a claim about evidence. When the skeptic argues that we fail to know that the zebra is not a cleverly disguised mule, she is making that claim that we do not *have any evidence for that belief*. The kind of evidence we have for the belief that x is a zebra usually does not involve the kind of evidence one would have for the belief that x is not a cleverly disguised mule.

1.5: Outline of remaining chapters:

In the remaining chapters I will argue that we can justify an assumption that sensory experience is reliable without appeal to an overly technical argument. Everyday knowers assume that sensory experience is reliable and I will argue that such an

¹⁸ Gail Stine, “Skepticism, Relevant Alternatives, and Deductive Closure,” from *Skepticism: A Contemporary Reader*, ed. Keith DeRose et al (Oxford: Oxford University Press, 1999).

assumption is epistemically rational because it is truth-conducive in important ways. The assumption that sensory experience is generally reliable, later referred to as AR, is rational in the sense that it is epistemically sensible in light of all of the information we have available to us. Yet, this is not to claim that we have *epistemic justification* for AR. Epistemic justification would amount to reasons or evidence that *counts toward* the truth of AR.

If we can show that the endorsement of AR need only be rational, then it is possible to highlight an anti-skeptical position that concedes to the skeptic that we fail to possess something as strong as epistemic justification in favor of AR while at the same time maintaining that our knowledge claims are justified. The argument that rationality is the appropriate requirement will rest on the notion that AR functions more like a presupposition than an explicit premise in the defense of our knowledge claims. Thus, I will conclude that since AR is not directly justification-conferring, it is only necessary that our endorsement of AR be rational. The notion of rationality here is such that it allows pragmatic and value oriented reasons to count as justification, whereas epistemic justification has traditionally been that kind of justification that necessarily speaks to the actual truth of the claim involved. Additionally, the requirement of rationality is weaker than that of epistemic justification. Epistemic justification requires reasons or evidence that speaks to the truth of the claim involved. Alternately rationality, while allowing for other kinds of justification, only requires that there be no evidence or reasons *against* the claim involved. Establishing the claim that an agent's reliance upon her sensory experience and her rejection of skeptical hypotheses need only be rational is the first step in an anti-skeptical position that admits the philosophical strength of skeptical arguments

while maintaining that the epistemic practices of everyday knowers are nonetheless epistemically justified.

The above somewhat subtle position seems most promising against the various realities that a truly coherent anti-skeptical position seems to face. Thus, I begin this discussion admitting that skeptics seem to be right about the fact that we fail to have compelling epistemic justification for our wholesale endorsement of sensory experience. They also seem to be right that *if* we had *epistemic* justification for our assumption about the reliability of the senses in general, then we would have epistemic evidence that excludes skeptical hypotheses. Of course, notoriously, we are unable to produce such evidence to silence the skeptic once and for all. The challenge of facing skepticism despite these admissions is the challenge of trying to make a fallibilist position about knowledge coherent despite a long tradition against fallibilism about knowledge.

I will begin by asking how it is possible that we might be justified in our external world knowledge claims while at the same time admitting that skeptical hypotheses may in fact be actual. Let me clarify that we are not aiming to defend the very counter-intuitive claim that it makes sense for one to say in one breath I know that I have hands, while in the next admitting that it may be true that I am a brain in a vat. Rather, we are seeking to find reasons to dismiss skeptical hypotheses that do not rest upon having evidence or reasons for thinking they are false; such evidence would amount to epistemic justification that defends our reliance on sensory experience. Instead, we want reasons to reject skeptical hypotheses that are sound, yet, do not ground knowledge claims about their falsity. The end of our argument will not be that we know that skeptical hypotheses are false. Rather, the end of our argument will be that we have good reasons for relying

on sensory experience *whether or not skeptical hypotheses are true*. Such reasons, admittedly, do not amount to epistemic justification. But if our assumptions about sensory experience are not justification conferring, it is not clear that the skeptic has any reason to require epistemic justification. Indeed, it does not seem that we operate with such standards in everyday life with regard to assumptions that are not directly justification conferring.

After articulating why rationality seems the more appropriate standard for AR, then we will look at whether or not we meet this standard. Namely, it must be shown that we are in fact rational to rely on sensory experience in a general sense. If we are rational to endorse sensory experience in general, this will imply that we are rational to ignore skeptical hypotheses. Thus, this aspect of our argument will focus on grounding the claim that the everyday practice of treating skeptical hypotheses as irrelevant is in fact a justified and sound epistemic practice. Again even if we are rational to ignore skeptical hypotheses, it will not follow that we know they are false. Rather, it will simply follow that we are rational to ignore them unless and until there is epistemic justification that indicates such a hypothesis is in fact likely. Thus, if it were to turn out that at some point we stumble upon evidence that directly indicates that we are in fact brains in vats, that it is no longer a mere possibility, then we would no longer be rational to ignore such a hypothesis.

In the next chapter, we will look at Laurence Bonjour's and Duncan Pritchard's work, as well as others, to give a sense of the specific way in which we are approaching the skeptical problem. This approach will involve focusing on the subjective questions that skepticism raises: questions such as, am I rational to rely on sensory experience, are

any of us ever justified in claiming knowledge of the external world, etc. Our brief discussion of Pritchard's work should lend some clarity as to why the focus of our discussion will be knowledge ascriptions when the skeptic thinks of herself as attacking knowledge rather than *merely* attacking *claims* to know.

The next step will be to outline in more detail the ways in which we might defend mere rationality as the appropriate standard regarding our reliance on sensory experience. The *rationality* criterion is to be defended against the skeptic's more traditional requirement that in order to be justified in our knowledge claims we must possess *epistemic justification* for our reliance on sensory experience.

After defending the notion that rationality is sufficient in order for us to count as justified in our claims to know things about the external world, we will explore the idea that we are in fact rational to rely on sensory experience. Since it is admitted that we do not have epistemic reasons to rule out skeptical hypotheses, this claim may be formulated as the claim that we are rational to ignore skeptical hypotheses even though it is logically possible that they are true. We will perform a thought experiment formulated initially by Henderson and Horgan in order to show that we are rational to rely on sensory experience even though we do not definitively know that sensory experience is in fact reliable. Prior to the thought experiment, we will look at the pursuit of truth in the context of everyday life in order to begin to articulate the sense in which we are rational to rely on sensory experience even though we do not know from our first person perspective which world is actual.

It has been said of the skepticism debate that one man's modus ponens is another's modus tollens, or that skeptics embody the pessimist's position, and non-

skeptics, the optimist's point of view. There is a sense in which the discussion to follow will acknowledge in some ways a kind of stalemate with the skeptic. Rather than pursue a traditional solution to skepticism that involves showing some way in which our ordinary beliefs are more likely to be true than skeptical alternatives, I am aiming specifically to undermine the skeptical challenge that our endorsement of sensory experience is *irrational*.

A concern of irrationality arises in the face of skepticism if we reason in the following way: beliefs maintained on sufficient epistemic grounding are rational, and those maintained in light of an awareness of *insufficient* grounds are *irrational*. Next, the skeptic might argue that given her skeptical argument, we have been made aware of how our external world beliefs are insufficiently grounded. Our claims are, according to her, insufficiently grounded because the assumption upon which they rest, the assumption that sensory experience is generally reliable and that we are not brains in vats, has not a shred of epistemic evidence in its favor. Thus, she might argue that our knowledge claims are not sufficiently grounded and that we ought to abandon such claims. She might argue further that if we do not abandon such knowledge claims in light of their poor epistemic standing, then we are irrational believers.

To summarize I aim to highlight an anti-skeptical position that concedes to the skeptic that we fail to possess epistemic justification for our reliance on sensory experience as opposed to her skeptical hypotheses. This reliance is tantamount to operating with the assumption that we are not brains in vats and that we in fact exist in a normal world. The skeptical challenge, as conceived here, is to understand how we might be justified in operating as though we live in a normal world while acknowledging

that skeptical hypotheses might be true. This is not to say that I aim to defend the idea that knowledge claims are consistent with such an admission. Rather, I aim to articulate how we might be justified in rejecting and ignoring skeptical hypotheses on grounds that do not amount to showing they are less likely to be true than their non-skeptical counterparts.

The second part of the strategy highlighted above is to show that we are in fact rational to endorse sensory experience. In other words, we are rational to behave as though we are in a normal world even though we might in fact be brains in vats. It will be argued that we are rational to behave in such a way in the sense that doing so allows us more and better epistemic goods than to behave in a skeptical fashion. Thus, the aim of this argument is to defend the optimist against the pessimist's claim that she is irrational when she thinks she knows such basic things as that she has hands, or that there are trees in her front yard, or that she sees a cat in her living room, etc.

Chapter 2.1: Laurence Bonjour and Epistemic Perspectives:

In *Epistemic Justification*, Laurence Bonjour argues in favor of internalist foundationalism as the center of his anti-skeptical position.¹⁹ He favors an account of justification in which justifiers must be accessible via introspection, and argues justifiers grounded in sensory experience play a foundational role. The foundational nature of perceptual experience stems for Bonjour from the direct and unmediated role that perceptual experience plays with regard to external world beliefs. While the bulk of those arguments do not interest us here, some of what he develops regarding the notion of epistemic perspectives is of interest here. In this chapter, we will explore and build upon Bonjour’s notion of epistemic perspectives and take a look at what he took to be the relationship between the first person perspective and skepticism. We will be interested in Bonjour’s distinction between first and third person perspectives to both illustrate what aspect of the skeptical argument is of import here, and to highlight how many everyday epistemic practices differ from those found in the philosophy classroom. The distinctions between the two contexts reach beyond the mere play of pragmatic versus epistemic concerns. It will be argued that the context of everyday life yields epistemic concerns distinct from those that may be found in more philosophical contexts. The distinctions between these two contexts are not precise, or without overlap, but they are significant enough to play a role in justifying the practices of everyday knowers.

¹⁹ See Bonjour and Sosa, *Epistemic Justification*, 1.

For our purposes here, we are interested in Bonjour's idea of first person versus third person perspectives on epistemic issues. The distinction between these perspectives stems from the nature of epistemic inquiry in each. In the first person perspective what concerns us in epistemic inquiry is the de facto status of our beliefs. Thus, I might be worried about whether or not *my* external world beliefs are truly justified, or may worry as to whether *our* external world beliefs are justified. In contrast, the third person perspective generally involves concerns about the nature of justification in general, or the nature of knowledge in general. The third person approaches questions of knowledge and justification from the perspective of an outsider looking in. Bonjour characterizes the first person perspective in epistemology and its' relationship to skepticism as follows:

In contrast, the internalist approach becomes essential, I suggest, when the issue is, not the third-person question of whether someone else's beliefs are true or reliably arrived at, but instead the first-person (singular or plural) question about the truth (or reliability) of my own or our own beliefs, especially the relatively global version of this question in which it is all of a person's beliefs that are in question. In relation to this global question, no externalist approach is available to the person or persons in question without begging the very question at issue (though some other person or persons could, of course, conduct such an investigation from the outside).²⁰

As we see here, Bonjour characterizes the first person as the stance we take when we are concerned about determining the *status* of our beliefs.²¹ Thus, we might anticipate that the first person perspective relies heavily on justifications available to us via introspection. Additionally, we might anticipate the first person perspective to involve

²⁰ See Bonjour and Sosa, *Epistemic Justification*, 37.

²¹ It should be noted that others have made very similar distinctions. Richard Foley talks about what he calls an "ego-centric" rationality, and as mentioned in the last section of this chapter, Audi talks about the process of justifying versus the state of being justified. These distinctions are very similar to Bonjour's notion of first and third person justifications.

many *norms* governing belief formation and justification. This is likely due to the fact that the first person perspective is the perspective we take when we are judging our beliefs and the beliefs of others and attempting to determine which beliefs are justified, and amount to knowledge, and which beliefs do not.

A third person perspective, on the other hand, approaches knowledge and justification from the perspective of an outsider. Thus, we might anticipate a third person perspective to focus largely on the following kinds of questions: what does it take for an agent to possess knowledge, what makes a belief (in general) justified, and what kinds of methods of belief formation are reliable? When we are in the perspective of an outsider, we can set the criteria for justified belief and the possession of knowledge, without addressing or answering questions about the de facto status of *our own* beliefs.

A third person perspective, according to Bonjour, involves “employing the various methods from the outside and assessing their success from that perspective”.²² So a third person approach to knowledge would look at the various methods used in belief formation and assess which methods are most successful. Consequently such an approach to knowledge will involve criteria for the possession of knowledge that are likely to be external to the agents’ conscious or reflective awareness. For example, from the perspective of an outsider it might seem plausible to suggest that if our sensory experience is in fact reliable, if we are hooked up in the right way to our world, then we often gain knowledge from our sensory experience. However, it is a very different and much more difficult question to figure whether or not my sensory experience is actually reliable.

²² See Bonjour and Sosa, *Epistemic Justification*, 37.

Because the first person perspective involves asking questions about our own beliefs, answering questions in this perspective will often involve an appeal to what is available to us via introspection. Therefore, Bonjour argues, internalist approaches to justification are relevant when our concerns are in the first person, and an externalist approach to justification is helpful when our questions are formed in a third-person way. However, it seems plausible that the internalist is posed to answer the general questions of the third person perspective as much as the externalist. It is less clear that an externalist account is helpful when we want to know whether or not we in fact possess knowledge, i.e. we seek to know whether or not we know. Such high order beliefs involve a reflective process, and require transparency of the epistemic status of our beliefs in a way that is likely involve distinctly internalist type justifications.

Before going on, it is worth taking another look at the issue of whether or not the first person maps onto an internalist approach to justification, and how well the third person perspective maps onto externalism. While it seems true that these two perspectives fit the internalism/externalism distinction, we should note that it is not a perfect fit. For example, it seems possible, if not likely, that we may be concerned in the first person, for example, about evidence that is not immediately available to us via introspection, and, therefore, requires some investigation. Such instances may undermine the sense that the first person perspective always takes an internalist approach to justification. For these reasons, I will largely avoid claiming that the first person perspective maps onto internalist type justifications. Instead, it seems important to follow Bonjour here and stick with the weaker claim that the first person perspective *lends itself*

to an internalist approach to justification, and avoid thinking of this relationship as stronger than that.

It seems the same can be said for trying to map the third person perspective onto the notion of an externalist approach to justification. However, perhaps this should not surprise us. These perspectives, while distinct epistemic concerns, and distinct approaches to the questions of justification and knowledge, are not entirely separate. It seems somewhat obvious that we move between these perspectives and can bear in mind concerns and interests from one perspective when we are in the other.

Such a blurring between these perspectives seems to raise the question of whether or not the distinction between these perspectives holds clearly enough to be of interest. My own sense is that interest in the distinction between these epistemic perspectives lies in the *explanatory power* of the distinction, and less in the notion that these perspectives are perfectly clear and distinct. The distinction allows for an *explanation* of the role and importance of both internalist and externalist approaches to justification. As Bonjour argues, both perspectives are important and each lends itself better to either an internal or external approach to justification. Building upon Bonjour's notion of epistemic perspectives, we may also see a way to explain such "epistemic phenomena" as Wittgenstein's hinge propositions, and epistemic issues with deductive closure. And, as I will attempt in the following chapters, to use this notion of epistemic perspectives to explain why the skeptic goes awry by arguing that some alternatives are not relevant. Prior to moving on to these arguments though, let us return to developing the notion of epistemic perspectives as articulated by Bonjour.

Bonjour argues that an internalist approach is essential in addressing the primary concern stemming from skepticism about the status of our beliefs. Such an internalist approach is required because an externalist approach in this instance will inevitably beg the question. Since a third person approach involves looking at belief formation and justification from the outside, such assessment over the success of an agent in forming beliefs will involve the *very methods* at question in skeptical arguments. In other words, any third person assessment of the reliability of the senses will involve the senses themselves. Therefore, Bonjour argues, the third person perspective is of limited use to us if we are concerned with skepticism as it pertains to the issue of the status of our own beliefs about the world.

Of course, we will note that a third person approach to skepticism can be of tremendous help, as many externalists have argued, if our concern in skepticism is to block the skeptic's anti-knowledge conclusion.²³ Such a third person approach to skepticism may allow us to block the skeptic's conclusion by arguing that one need not *possess a reason* for thinking that the senses are reliable in order to have knowledge via sensory experience. Such a move is clearly grounded in a third person perspective because it involves spelling out conditions for the possession of knowledge versus claiming any particular beliefs *do in fact* amount to knowledge. Since we might see the skeptic's argument as formulated on the grounds that we need a reason to think that sensory experience is reliable in order to have a justification for preferring our ordinary beliefs over their skeptical counter-parts, we may avoid skepticism by taking such an externalist approach to justification and knowledge. Yet, we might note along with

²³ For an example of such externalist type moves, see Ernest Sosa's argument in *Epistemic Justification*, John Greco's *Putting Skeptics in their Place*, as well as Timothy Williams' *Knowledge and Justification*, and Duncan Pritchard's *Epistemic Luck*.

Bonjour, that such a response really does not do much to squelch the worry we might develop about the status of *our own* beliefs in light of skepticism.

The externalist response to skepticism gives us a conditional solution to skepticism. *If* our sensory experience is reliable, then our true beliefs formed on the basis of sensory experience likely amount to knowledge. Yet, we cannot confirm in a non-circular, non question-begging way, whether or not our senses are in fact reliable. Nor does the externalist articulate reasons that would justify *our* reliance on sensory experience. As the skeptic notes, we have no non-question begging reason to think sensory experience is reliable *in general*. Thus, solutions to skepticism grounded in such third person moves as those outlined above, may well leave us with worries and concerns revolving around the first person question about the status *of our beliefs*. Following Bonjour, we will focus on this distinctly first person worry about the status of our beliefs.

To summarize, it is highly likely that any third person or externalist attempts to show the reliability of the senses is likely to involve sensory experience itself. A third person approach might be able to stipulate that knowledge is attained when reliable methods are used, yet it cannot show which methods are in fact reliable without using the very methods in question.

So, along with Bonjour, we might conclude that an externalist approach to skepticism is of some help from the fact that externalists do not require agents to possess an internalist justification for their beliefs. Thus, we might expect the following from an externalist approach to skepticism: if sensory experience is in fact reliable, then we possess knowledge of the external world whether or not we have any justification for *thinking* that sensory experience is reliable. Such an answer might seem to offer some

relief in the sense that we no longer need *to find* some evidence that favors our beliefs about the world over skeptical hypotheses. Such has often been thought to be the crux of skepticism. However, Bonjour rightly concludes that such an account does not offer us non-question begging reasons for thinking sensory experience is in fact reliable, nor does it seem able, in principle, to do so. Bonjour claims that one of the primary epistemic concerns stemming from skepticism is the first-person question of whether or not *my beliefs* are justified. This concern is distinct from the question of what it takes to possess knowledge or have justified beliefs – I want to know whether or not *I* possess knowledge or whether or not *my* beliefs are justified. To answer *that* question, I must look to the reasons *I* have for my beliefs.

Bonjour and others who have argued similarly seem to be right about this issue. Even if an Externalist approach to knowledge addresses some of the skeptical problem and seems to win the day as an approach to knowledge, there are deep epistemic issues in skepticism that remain in light of such purported solutions. I will look at this issue a little further in chapter 3, however, for now, we can see the reasons for approaching skepticism in this more first-person fashion. Thus, it is the first person question of whether or not our beliefs about the external world are justified that we will focus on in the discussion that follows.

2.2 Epistemic Perspectives and Skepticism

Let's turn now to the issue of diagnosing skepticism. We should remember that I have chosen to focus on the under-determination characterization of skepticism. I have

focused on this characterization of skepticism because it seems to be the core of skepticism and it is the most pernicious version of skepticism.

On this characterization of skepticism, the skeptic argues that our ordinary beliefs about the world do not amount to knowledge because we lack evidence that favors our ordinary beliefs over other various skeptical hypotheses. Since we have no reason for preferring our ordinary beliefs over skeptical hypotheses, we fail to have sufficient evidence and fail to count as knowing. Alternately, you might give a slightly different version in which a skeptic argues that we fail to count as knowing our ordinary beliefs because we fail to possess any (non-question begging) evidence that our sensory experience is reliable.

What I will argue is that the skeptic is pointing to a more significant and general epistemic problem. If we keep in mind Bonjour's distinction of the first and third person I think we can assess external world skepticism as a version of a more general epistemic issue. The more general epistemic issue is the gap between first person, internalist type justifications and externalist requirements for knowledge. Such a gap has long been recognized. Indeed, philosophers have addressed this issue in the form of articulating how knowledge is possible in the face of our intrinsic fallibility, the fallibility of the evidence upon which we base our beliefs, and the problem of the criterion.

Alternately, philosophers have wrestled this issue in the problem of Gettier -type examples. In such examples an agent possesses a belief that is both justified and true and yet fails to count as knowledge because, unbeknown to the agent, their evidence does not connect up to the truth in the right way. A well worn example is of an agent who looks at the clock in her office and notes that it says 4 o'clock, and she forms the belief on that

basis that it is 4 o'clock. As it turns out, unbeknownst to the agent, the clock has a dead battery and is not working. It also turns out that her belief is true – it just so happens to be 4 o'clock. The intuition in such an example is that her belief was justified and it was true, but she did not possess knowledge because of the accidental nature of the connection between her justification and her belief being true. It turns out that we tend to think that justification must be connected to truth in such a way that excludes knowledge as the result of a lucky guess.

Indeed, we can see Gettier cases as showing that the gap between justification and truth is greater and more epistemically problematic than previously thought; so much so that the traditional account of knowledge as justified true belief seems insufficient. I will argue that the best way to understand skepticism is fitting into this larger epistemic issue. My aim is to attempt to articulate this issue in a clear way and to address what it might look like to attempt to solve this issue.²⁴ In the context of first person justifications, Gettier concerns and skeptical worries in general are prevalent and pernicious. As suggested though, this ought not to count against internalist type justification, since it plays such a pivotal role in our first person epistemic endeavors. Regardless of whether or not one might embrace externalism as an account of knowledge, it remains that we are interested in assessing the actual epistemic status of our beliefs. And we only have what is available to us via introspection to make such judgments.

Before moving on, let's spend a moment talking about the first-person perspective and skepticism. One very pertinent way to see skepticism involves looking at it as a

²⁴ Additionally, though I don't want to argue for this, it is my sense that epistemologists have addressed this issue in a fragmented and insufficient way – in the sense that this epistemic problem has not been spelled out or brought to the surface in a way that allows it to be addressed head on.

challenge posed to our beliefs and our evidence by an imagined “skeptic”. This way of thinking of skepticism makes the challenge that the skeptic puts forward against our beliefs very immediate and personal. The challenge put forward by the skeptic is one against our own beliefs. Thus skepticism seems, for many, to cause a first-person worry about the status of one’s own beliefs. If we unpack this worry a bit, we feel a charge of irrationalism leveled against us by the skeptic. According to her, we believe without sufficient grounds for believing. We buy into the reliability of the senses without any evidence that epistemically justifies such faith. This latent charge of irrationality embodied in the skeptic’s argument is the focus of the discussion to come.

Characterized in this way, we can understand why skepticism has lured so many thinkers to engage this particular paradox about knowledge. Additionally, it seems right to think that this first person worry is one of the most potent worries arising from skepticism.

What I would like to suggest here and spell out in more detail later, is that skepticism actually arises out of our ability to shift from the first to the third person perspective combined with the fact of fallibility for first person evidence and reasons. In chapter 4 I will attempt to fill in some of the detail of this view. For now though I would like to look at how we might characterize skepticism as rising out of shifting epistemic perspectives and fallibility. Admittedly though such a conception is rather vague and general and I will ultimately only be able to outline such a view. However, since it seems to my mind to be a part of the larger picture in defending a fallibilist view of knowledge, it seems worth hinting at or bearing in mind even if working out the details is beyond the scope of this work. The bulk of the argument here is that the assumption that we are

agents in a normal world and that sensory experience is reliable need only be a rational assumption. The second portion of the primary argument is that assuming the reliability of the senses is in fact rational. However, we might bear in mind that part of the argument that rationality is the appropriate requirement will implicitly involve the notion that first person justifications are limited in the sense that they do not definitively connect us with truth. Since they do not so connect to truth, it seems appropriate that the primary assumptions involved in first-person justifications need only be rational.

In order to see how skepticism might arise as a result of shifting perspectives we might consider how an imagined skeptic might come to worry about skepticism, and how ordinary believers might also come to see skeptical worries. First though let us state the view briefly. In a third person way, when we look from the perspective of an outsider at what it takes to know, we feel that justification is connected to truth. It must be that justifiers indicate the truth. It is this connection between the justifier and truth that prevents a lucky guess from counting as knowledge. Thus, our JTB understanding of knowledge has always implicitly endorsed the view that the justification must be connected to the truth of the belief. Similarly, in a good argument, the truth of the premises must support the truth of the conclusion; they must have the right relationship to one another. However, when we move into our first person perspective, what we have access to via introspection is not sufficiently connected to truth, as the skeptic is happy to point out. The skeptical challenge is really a challenge about resolving the standards for knowledge set from the perspective of an outsider with what we actually have available to us through introspection. If we are to solve the skeptical challenge we must conceive of what we have access to in the first person as consistent with the possession of

knowledge. This challenge is likely a challenge of conceiving of fallible knowledge at the theoretical, third person, level. What the skeptic does is take her conception of what it takes to know from the perspective of an outsider and concludes that we do not meet those criteria in the first person – that is, we fail to meet the standards of adequate justification. However, it is possible to do the opposite. Rather than take the standards set from the third person as definitive, we might begin with our first person intuition that we do in fact know many things about the external world, and ask ourselves what must be changed about our theoretical conception of knowledge in order for what we possess in the first person to count as adequate in the third person perspective.

Prior to moving on, we should note that characterizing skepticism as an epistemic worry or crisis that results from shifting epistemic perspectives allows us to *explain why* skepticism seems to be so difficult to maintain as a positive belief system in the course of everyday life. If skepticism really results from looking at knowledge in a third person way and then applying that to our first person perspective, we should expect skepticism to be difficult to maintain or practice in the first person. So the skeptic is a skeptic when looking at her beliefs about the world from a third person point of view but she struggles to maintain her skepticism when she shifts more firmly into her first person point of view (in which she is again processing the information at hand). She quits viewing sensory experience with suspicion as soon as she shifts back into a first person perspective and goes about her daily life. It is only the perspective of an outsider that tells her she is not warranted in relying on sensory experience, yet her first person drive continues to be that of pursuing truth and a strong sense that sensory experience is the only avenue for doing

such. Such conflict is more than psychologically interesting; it is philosophically interesting as well.

2.3: Nagel's View From Nowhere:

Nagel's conception of moving perspectives, and characterizing this move towards objectivity as a matter of taking "up a new, comprehensive viewpoint after stepping back and including our former perspective in what is to be understood" may be helpful to us here.²⁵ In theoretical reasoning, as we shift from a more self-focused perspective to one that incorporates ourselves as components but is neutral in the sense that it does not represent any one person's point of view, we may form a new set of beliefs. Bonjour's concept of epistemic perspectives is part of a larger aspect of reasoning in general. When we talk about the third person perspective in epistemology I think it may be helpful to have in mind Nagel's idea of a neutral or objective viewpoint. On the other hand, the first person perspective represents a given agent's particular epistemic viewpoint.

One way to think about Bonjour's claims is that we have distinct epistemic perspectives, and that these distinct perspectives result in different conceptions of justification. The resulting aspects of justification will, roughly, match the internalist and externalist approaches to knowledge. While Bonjour's primary concern is to see how skepticism is primarily a first person concern and, thus, is best handled by an internalist approach to justification, my interests lie in the epistemic relevance of these differing

²⁵ Thomas Nagel, "Value: The View from Nowhere," in *Ethical Theory*, ed. Louis Pojman (Wadsworth publishing 2007), 169.

perspectives and how these different perspectives might play a role in generating the global skeptical worry exhibited in skepticism.

What we might glean from Nagel is the sense that this third person and objective viewpoint is distinctly more theoretical than our subjective viewpoints. That is, the objective viewpoint more readily allows for systematic knowledge. For the purpose of our discussion to come I find Nagel's view of objective viewpoints and theoretical knowledge interesting because we may want to make an association between first person epistemic practices and everyday life, and third person practices and philosophical endeavors.

It need not be the case that the first person coincides perfectly with our everyday epistemic practices while third person questions coincide with the distinctly more philosophical. However, what I would like to suggest, in part, in the chapters to come is that skepticism is easier to accept in the third person and more philosophical contexts than it is in everyday life. This is so because the theoretical implications of skepticism are less immediate and less concrete than the first person implications. Thus, making it significantly easier to be a skeptic in the Philosophy classroom than it is as a pedestrian on the street.

The above claim is not as unusual as it might initially seem; both Descartes and Hume admitted as much in their skeptical writings. However, both attributed the difficulty of maintaining the skeptical attitude to psychological or instinctual influences rather than to legitimate epistemic ones. Thus, one of the projects of this paper is to outline a sense in which the difficulty in maintaining skepticism as a coherent practice lies in legitimate epistemic concerns rather than the merely psychological.

We might see this particular issue as the challenge to *explain* a certain bit of data. That bit of data is the overwhelming fact that there are very few practicing skeptics in the world. Many are non-skeptical in the sense that they operate on a day to day basis with the assumption that sensory experience is in fact reliable. As we well know, the skeptic feels this assumption is not warranted. The skeptical explanation for the lack of practicing skeptics is that we are psychologically compelled to believe in sensory experience, or that we are simply not critical enough in our assumptions. Whatever the explanation the skeptic might opt for, her explanation involves a denial of legitimate epistemic reasons for embracing sensory experience. In the discussion to follow I would like to suggest that the skeptic is mistaken when she leaps from the fact that we have no epistemic justification for relying on sensory experience to the conclusion that there are no epistemic reasons for relying on sensory experience.

As part of the argument to follow, I will argue that there are concerns both epistemic and pragmatic that are unique to the first person perspective. Since I will generally assume that it is safe to contend that this is the primary perspective we take in everyday life it will turn out on my view that skepticism is epistemically more difficult to embrace in everyday life than it is to embrace in the philosophy classroom. Thus, unlike Descartes and Hume my sense is that the difficulty with maintaining one's skeptical position after the philosophizing is done is that there are important epistemic distinctions between the two contexts. These distinctions are such that even those who might be sympathetic to skepticism in the philosophy classroom will find it an unwise position in the course of everyday life.

2.4: Audi and the *Structure of Justification*:

In his “Justification, Truth, and Reliability” Robert Audi talks about two quite different strands of justification. He characterizes these distinct strands of justification as deontological and ontological, or as the *process* of justifying versus the *property* of justified.

In this paper, Audi argues that it is important in epistemology that a theory *illuminates* the relationship between justification and truth; that it illuminates the relationship between the property justified and the process of justifying. Looking at Audi’s view in more detail, he argues that an ontological approach to justification will involve the claim that “when something justifies a belief, then, in a suitable range of relevantly similar possible worlds, notably worlds like ours where the same sorts of things is believed on the same sort of basis, this belief is true”.²⁶ In distinction, the teleological view of justification involves tying “justification to *seeking truth* and makes the *practice* of justification, above all the giving of justifications, fundamental in the connection between concepts of justification and truth”.²⁷

Roughly, we can say that these two aspects of justification will mirror the internalist and externalist distinctions in approaching justification and knowledge. At least in the under-determination version of skepticism, the skeptic seems to approach justification from the teleological perspective. Challenging others to provide evidence for their ordinary beliefs over skeptical hypotheses, the skeptic asks others to go through

²⁶ Robert Audi, “Justification, Truth and Reliability” in *The Structure of Justification*, (Cambridge University Press, 1993).

²⁷ See Audi, “Justification Truth and Reliability”, 301.

the process of justifying their beliefs, and, when they are unable to do so, she concludes that they fail to possess knowledge.

Audi contends that epistemologists have failed to take into account the relationship between the property of justification and the process of justifying. He talks about a need to articulate further the relationship between truth and justification. Audi proceeds to propose a number of principles for prima facie justification for different belief forming processes (e.g. perception, memory, etc.).

What is interesting for our purposes is that while Audi's distinctions and positions are similar to the view being developed here, I am more inclined, perhaps along with Bonjour, to see these two strands of justification as *stemming from* different epistemic perspectives. In Audi's language, I am exploring the idea that not only do we need to further explore the relationship between justification and truth, we need to further explore what epistemic *obligations* we may have in light of these different ways of thinking about and approaching justification. Skepticism exemplifies instances of shifting epistemic perspectives and the myriad ways in which core epistemic issues actually arise from shifting epistemic perspectives.

In the context of skepticism, it seems that the skeptic begins by focusing on *justifying* ordinary beliefs and shifts to an objective perspective in which she asks herself about justification as a property, as it relates to *truth*, and in that perspective she notes that the evidence that serves as justification for our ordinary beliefs fails to imply truth (she notes this via the consistency of our sensory experience with being a brain in a vat, for example). While I do not want to contend that the skeptic's shifting perspectives is necessarily problematic, I do think it raises the general and very deep epistemic question

about the connection between justification and truth. Further, it raises deep epistemic questions about shifting perspectives in general and what potential epistemic issues that may arise from shifting perspectives. Prior to looking at a solution to skepticism, we may note the way in which I acknowledge skepticism as raising important and deep questions in epistemology. The skeptic is asking precisely the questions Audi thinks are the deepest in epistemology; she challenges the connection between our justifications that serve in our processes of justifying with truth.

Skepticism raises the very deep epistemic question about how we are to handle justification in the first person perspective in light of the fact that we can see that justification fails to imply truth in the vast majority of cases. In a sense, we might even see skepticism as noting that many of our usual justifications fail to *imply* truth. And while we might be inclined to address skepticism simply by rejecting the skeptic's standard that justification ought to imply truth (i.e. that our evidence *ought* to speak against skeptical hypotheses), Audi is right to claim that we fail to possess an adequate conception of how justification *is* connected to truth, if it fails to *imply* it. It does seem, as Audi claims that it is "at least partly constitutive of justification that, in *some* way, it *counts toward truth*".²⁸

While we will not address the issue of the connection between justification and truth in any explicit way, one of the ways we may think about the distinction between epistemic perspectives is that it tends to involve different aspects of justification. Thus, one way to approach the question of justification and truth may be by outlining the

²⁸ Audi, *Structure of Justification*, 301.

relationship between a first and third person perspective. Unfortunately, there is only room enough in this discussion on skepticism to explore this issue briefly in chapter 4.

2.5: Duncan Pritchard's *Epistemic Luck*:

In his *Epistemic Luck* Duncan Pritchard characterizes the heart of skepticism as concern over the tension between the possession of knowledge and various kinds of epistemic luck. Epistemic luck in an extreme variety is that phenomenon of guessing and getting the right answer. While Pritchard notes several varieties of epistemic luck, the two kinds that pose the most problem for the possession of knowledge are what he calls veritic and reflective epistemic luck. Veritic epistemic luck is that kind of luck such that it is a matter of luck that the agent's belief is true.²⁹ Veritic epistemic luck is the kind involved in Gettier style examples. In the clock example we talked about previously, it was enormously lucky that the agent's belief that it was 4 o'clock was true. It was enormously lucky in the sense that in most possible worlds, a belief formed on the basis of a broken clock would not yield a true belief about what the time is. Pritchard contends that while an externalist account of knowledge that embraces some version of a safety principle will address or eliminate veritic epistemic luck, reflective epistemic luck remains in light of such views, is difficult to remove, and makes a prima facie case for an internalist epistemology. Safety based externalist positions eliminate veritic luck as consistent with knowledge because they require that a belief must be safe if it is to be knowledge. Pritchard develops roughly the following safety principle: that in all nearby worlds where the agent forms her belief in the same way as the actual world, she only

²⁹ Duncan Pritchard, *Epistemic Luck* (Oxford: Clarendon Press, 2005), 146.

believes when the belief in question is true.³⁰ It is clear that Gettier cases will fail to cases of knowledge on such an account. Such a safety based view does not, however, eliminate reflective epistemic luck.

Reflective Epistemic Luck is that brand of luck in which from the agent's own perspective it is a matter of luck that her belief is true.³¹ It seems as though this is the kind of luck that is involved in skeptical arguments and is not helped by a safety based account of knowledge. Pritchard states of reflective luck that "Significantly, this type of luck can remain even if we stipulate that the agent in question has a true belief that is safe."³² Thus, it might be the case that in most nearby worlds the agent only believes in those instances in which her belief is true, and, yet, from her own perspective it is a matter of luck that her belief is true. Thus, we might imagine an agent believes that she has hands. We might also imagine that in most nearby worlds she only believes that she has hands in those instances in which it is true that she has hands. Thus, we might imagine that her belief is in fact true because she is in a normal world, and her belief is safe. However, it seems to remain that from *her own perspective* it is a matter of luck as to whether or not her belief is true because she does not know that she is in a normal world (as opposed to a brain in the vat world).

Thus, in the above example reflective luck remains and seems inconsistent with the agent claiming that she knows she has hands. How can she know she has hands if for all she knows she doesn't have hands and is a brain in a vat? This is the heart of the first person worry that skepticism raises. Skepticism narrows in on the issue of reflective

³⁰ Pritchard, Epistemic Luck, 163.

³¹ Pritchard, Epistemic Luck, 173.

³² Pritchard Epistemic Luck, 173.

epistemic luck. Further the skeptic feels that in order to rule out the sense of luck we must have evidence that allows us to rule out skeptical hypotheses.

Both skeptics and internalists are worried about epistemic luck to such an extent that they require “full cognitive responsibility” in order for an agent to count as knowing. An agent is fully cognitively responsible when her belief is immune from even reflective epistemic luck. What skeptical arguments do is note that agents are not and cannot be immune in this way. In other words, the skeptic notes that despite the internalist’s requirements on knowledge, our beliefs are subject to a certain amount of reflective epistemic luck.

Pritchard characterizes the problem posed for internalist justifications by skepticism as an issue that occurs at the level of knowledge ascription. Thus, an externalist account of knowledge will save knowledge, but we cannot ascribe it to ourselves or assert it of anyone else in a skeptical context. Skeptical arguments undermine the epistemic practices occurring in the first person. Any purposive, conscious and deliberate epistemic projects we undertake are particularly affected by skepticism. As it stands, the skeptic has undermined these practices by undermining the epistemic value of the internalist type justifications that we use to carry out all of these practices. The skeptic has shown that sensory experience, for example, is not necessarily indicative of that which we take it to be (i.e. of a world full of three dimensional objects). This problem posed by the consistency of internalist type justifications (sensory experience in particular) with various skeptical possible worlds impacts all of the epistemic practices that use these kinds of justifications to navigate and assess our epistemic situation from the first person perspective.

The only way we have of saving the integrity of these practices is to resist the skeptic's assumption that internalist justifications need to speak against skeptical hypotheses in order to have epistemic value enough that we can rely on them for evidence. This is difficult, however, because it just is the case that our usual evidence is consistent with many skeptical hypotheses. So we cannot hope to avoid the skeptic's conclusion by claiming that we have reason for rejecting skeptical hypotheses in the sense of having evidence that favors our ordinary beliefs. We do not. Indeed, it seems that we have no more reason for rejecting skeptical hypotheses than the skeptic does for putting them forward in the first place.

Pritchard argues that skepticism essentially attacks a kind of “fine grained” internalist knowledge. Since the skeptic launches a successful attack on this kind of knowledge, all of our practices that involve the first person are damaged by the skeptic's argument. Duncan characterizes the skeptic's argument as pointing out how our beliefs are subject to a certain amount of reflective epistemic luck. This, he argues, does not imply that we fail to possess knowledge. That the skeptic has undermined internalist justification is something to take seriously though. Duncan Pritchard winds up with the position that there is a certain epistemic angst that arises when we reflect on our epistemic position. As he puts it, “Our epistemic position is thus akin to a high-wire acrobat who is unable to be sure that the safety net has been erected below. For sure, it *looks* as if it has, but it would look that way even if it hadn't been erected”.³³

Pritchard's characterization of our epistemic situation from the first person perspective strikes me as quite on target. What I would like to do in the rest of this chapter is explore in just a little more detail what exactly the concerns are for internalist

³³ Pritchard, *Epistemic Luck*, 248.

type justification given skepticism. I will urge that while Pritchard's account is accurate, we should not surrender to the skeptic entirely on the issue of knowledge ascriptions.

Pritchard uses Wittgenstein's hinge propositions as a way of accounting for why it is in the course of our everyday epistemic practices that certain propositions do not admit of evidence. Certain propositions act as hinge propositions from which the entire practice of giving and asking for reasons presupposes and depends upon. The evidence we possess for our everyday ordinary beliefs about the world does not support these hinge propositions. Hinge propositions are characterized as propositions that are held without grounding.³⁴ The entire practice of giving and assessing reasons depend upon these assumptions holding, if we doubt them then we cease to play the game of giving and taking reasons at all.

What this account is lacking is an explanation of why the hinge propositions, propositions that serve to justify in a given context but are not themselves justified, are hinge propositions. The account claims that they are necessary for the entire enterprise and must be presumed in order to even begin; however, we might be to say a little more than this. One of the things we will look for in the next chapter is a little more substantive explanation of why these propositions are indeed hinge. We will largely agree with the notion that general assumptions of reliability about sensory experience are generally required in order to make sense of the give and take of justifications regarding sensory experience in the first place. This notion may well ground the idea that AR acts a presupposition acts in an argument much more than it acts as an explicit premise.

Instead of conferring justification as premises confer justification on a conclusion, general assumptions of reliability function more like presuppositions in an argument.

³⁴ Pritchard, *Epistemic Luck*, 226.

Once the general assumption is accepted, sensory experience gains a foundational and direct evidential status on its own that is not epistemically dependent on the grounds for the assumption itself. Thus, the role of assuming the reliability of sensory experience is not a justification conferring role. Since our general assumption is not justification conferring, it is not necessary that we possess something as strong as epistemic justification in favor of it. Rather, it is sufficient for the assumption to simply be a *rational* assumption.

Thus, this is the starting point of our inquiry into responses to this kind of skeptical worry.

2.6: Summary

In light of our discussion thus far, we can characterize the first person perspective along with Bonjour as the perspective that we are particularly concerned about here. Alternately, the third person perspective, according to Bonjour, is the perspective in which we ask from an outsider's perspective what it takes to possess justification or knowledge. As we gleaned from Nagel, we can think of the first person perspective as involving our individual perspectives to a greater degree, whereas the third person perspective can be characterized as a more neutral perspective. Additionally, we looked at other ways to potentially characterize the notion of epistemic perspectives. To this end, we looked at Audi's discussion of the *aspects* of justification and the distinction between focusing on justification as a *property* versus *the state* of being justified.

There are still a couple of key distinctions that need to be made about the first and third person perspectives. For our discussion in the next chapter, it will be important to

bear in mind some concrete ways in which the first person perspective is distinct from the third person perspective.

Next we want to note a number of things about the first person perspective. In the first person perspective we are actively engaged in assessing our own beliefs, assessing evidence presented in a specific context, or judging methods of belief formation, etc. In this context, and within this perspective, we are actively forming beliefs as well as attempting to determine the status of our beliefs. Thus the practical pressures of forming beliefs in the context of going about our everyday lives are more apparent in the first person perspective.

It is in the first person perspective that we can see how everyday knowers may fail to be epistemically ideal. Everyday knowers are forming beliefs and assessing evidence in contexts with practical and pragmatic concerns that compete with and constrict the pursuit of truth. Indeed, when focused on the first person perspective, we must recognize other epistemic values than the goal of truth. It is also an epistemic value that we form beliefs in a way that meets the concrete and practical concerns of the context. So we seek not only to form true beliefs, but also to form beliefs that allow us to *successfully navigate our environments*. So even though this is a rough way of characterizing the first and third person and one that requires some qualifications, we may think of the first person as distinctly more everyday life and the third person as a distinctly philosophical perspective.³⁵

³⁵ Though I do not mean “philosophical” in an academic sense; the third person is philosophical in the sense that it is a perspective in which we contemplate the nature of fundamental concepts – what is it to possess justification, what is it to possess knowledge, etc.

2.7: Strategy for Handling the Skeptic:

This distinctly first person and, thus, largely internalist concern raised by skepticism pertains not only to the internalist but to any externalist for whom internalist justification plays a role in knowledge *ascriptions*. While many find an externalist approach to knowledge viable and appealing, it is less clear that an externalist account works equally well for knowledge ascriptions. To imagine that an agent could justifiably claim knowledge without any reasons for thinking she possesses knowledge is significantly more counter-intuitive than claiming that an agent can in fact know regardless of whether or not she has reason to think she knows. Thus, unless we are rabid externalists, a serious and significant problem remains for knowledge ascriptions even in light of an externalist solution to skepticism regarding *knowledge possession*.

Thru out the discussion to follow, I will focus on this specific worry about whether or not our *claims* to know are justified in light of our inability to rule out skeptical hypotheses. There may well be many issues of epistemological significance arising out of skepticism; we may be concerned to understand how knowledge is possible, we might wonder if and how knowledge is consistent with fallibility, or we might struggle to articulate a view of knowledge that is immune to skepticism. However, at the heart of the skeptic's criticism is the implicit claim that those who claim knowledge are somehow irrational, or in some way deficient in fulfilling their epistemic duties. Implicit in many skeptical arguments is the claim that the skeptic's position, the skeptic's reluctance to claim knowledge, reflects the *epistemically appropriate* position in light of our inability to rule out skeptical alternatives.

This last worry, the above criticism from the skeptic, is the challenge that is at the heart of the first person worry about the status of our knowledge claims. I am concerned in this discussion to address this specific question; I will not give a full treatment of all of the issues rising out of skepticism. Thus, our primary concern will be to address the worry that somehow our knowledge claims are not justified. If they are not justified then those who make such claims are somehow epistemically deficient; they claim knowledge when it is not appropriate, and, as the skeptic claims, they ought, as she does, to refrain from doing so.

This portion of the skeptic's challenge is going to boil down to a claim that we are not justified in endorsing or relying upon sensory experience as a method of belief formation. Specifically, the skeptic will argue that we have *no evidence* that sensory experience is reliable in general. One of the arguments I will give in a later chapter is that there are no clear reasons to accept the skeptic's requirement that we possess some evidence that sensory experience is generally reliable in order to be justified in relying upon it. While it is counter-intuitive to imagine knowledge claims as not requiring an internalist justification, there is little to persuade us that this internalist requirement should pertain for the entire chain of justifying reasons. Thus, if our grounds for our ordinary claims about the world are our various sensory experiences, it is not clear that we need to accept the view that in order to be justified in our knowledge claim we must also possess evidence accessible via reflection that justifies our grounds meant to justify the first order belief about the world.

That the skeptic's first person challenge boils down to whether or not we are justified in relying on sensory experience can be seen in the following reasoning. In

general we meet the skeptic's challenge that we have reasons we can produce in favor of our claim that we have hands. For example, I might cite my sensory experience of seeing and feeling hands as good reason to think I have hands, and I might ground a knowledge claim in such experience. The skeptic's challenge is to point out that our sensory experience may not be reliable in general. The challenge has shifted from a defense of our belief that we have hands to defending our belief that sensory experience is reliable; she will point out that for all I know I could be a brain in a vat. Since I have no reason to think that sensory experience is generally reliable, then my belief that I have hands is not well founded. Therefore, this particular debate in skepticism will turn on whether or not agents are entitled to assume that sensory experience is reliable, and this is how I will frame the skeptical problem about knowledge ascriptions in the discussion to follow.

If we allow this debate to be framed in the way the skeptic would like, then it is over before it has begun. If the skeptic is right that we must have actual evidence to indicate that sensory experience is generally reliable in order to be justified in relying upon it, then the skeptic has won this debate. However, it is not clear what might motivate us to accept this requirement. The primary argument to follow will be that we should not accept this requirement of the skeptic's. One could argue that such a requirement leads to an infinite regress, and, therefore, is in principle and theory a skeptical position. Such a position lands in skepticism in theory in the sense that it leads to skepticism prior to even taking account of our particular epistemic circumstances. Such a theory would land in skepticism for any finite and fallible knower.

Rather than push this tactic against the skeptic's requirement, my strategy will be to evaluate this requirement on its own merits. I will evaluate what such an approach in

epistemology looks like, what it would look like to practice this view, and evaluate what might motivate this view. Thus, we will be able to see and make explicit the implications and deficiencies in the skeptic's approach that are not so explicit in the skeptic's argument. In doing so, we should be able to find some strong motivation to reject the skeptical position.

However, we should not anticipate finding a refutation to the skeptic's position. This is not the strategy I have in mind nor do I find it a promising strategy for any who think it is the one we ought to take. The skeptic is ultimately articulating a standard for rationality, and, thereby, articulating standards for claiming knowledge. We cannot appeal to a definition of knowledge or to some other technicality to refute what is essentially a *value* claim. The skeptic is articulating standards that embody what she takes to be the norms for responsible belief formation and the claims that are appropriate in light of such standards.

If we look at the contemporary literature in epistemology and skepticism, it is clear that what many want is a refutation of skepticism. The responses from semantic externalism, contextualism, and externalism all offer ways to avoid the skeptic's conclusion; each, of course, has its strengths and weaknesses. None of these approaches, however, addresses the skeptic's most stringent criticism that claiming knowledge in light of uneliminated possibilities is not appropriate or rational. The contextualist admits that it is not appropriate to claim knowledge in a skeptical context, and the externalist leaves our claims to know as justified in a merely conditional manner. We are justified in claiming knowledge *if* we in fact know, and whether or not we know will be determined based on whether or not sensory experience is in fact reliable. This last concern does

little to guide us on what is appropriate in the first person. Is it appropriate to claim knowledge despite that we have no evidence to indicate that sensory experience is in fact generally reliable? This fundamental first person worry remains in light of these many contemporary responses to skepticism, and it is this issue that I will be concerned with in the following discussion of skepticism.

The solution to follow will essentially be grounded in claims that the skeptic is wrong about the values she articulates in the standards she endorses. The strategy is to make explicit what values the skeptic's standards embody and to evaluate those in a normative fashion. Thus, the argument to follow will defend the claim that agents who assume the reliability of the senses embody better epistemic values than the skeptic does. I will grant now that such an argument may not feel definitive for some, but I believe it is the kind of argument that is needed to address this particular first person worry stemming out of the skeptical argument. It is important to note that this worry stems from the normative elements involved in the first person perspective. We are particularly concerned with responsible and rational belief formation in the first person.

A worry might rise out of skepticism that our knowledge claims are not justified—that they are epistemically lacking or even irresponsible in light of the competing uneliminated error possibilities. It is this particular worry or charge I intend to evaluate in the discussion to follow. While I will not give a solution that provides us definitive grounds for ruling out skeptical alternatives, I will give a discussion that takes the sting out of the skeptic's particular charge that our knowledge claims are irrational in light of skeptical hypotheses.

A significant first step in the strategy I have outlined is to substantiate the claim that our reliance on sensory experience need only be rational. This step may well involve the most theoretically substantive aspect of the proposed strategy, and, yet, I will only be able to outline general reasons in favor of this position. I suspect that endorsing a looser requirement on assumptions of reliability than the skeptic has in mind is a move towards the fallibilist position that the skeptic does not embrace. Further it seems that such a move is a reflection of where one places the bulk of what has epistemic value. As in the case of many disagreements occurring at a fundamental level, it is unclear that this disagreement with the skeptic boils down to anything more than a massive conflict of intuitions. Nonetheless, I would like to offer some measure of justification for the position that assumptions of reliability need only be rational rather than supported by direct evidence of their truth which is the skeptic's position.

We have boiled the skeptical problem down to a problem about determining the nature and status of the reasons we have for relying on sensory experience. As the skeptic will point out, after all, we could be massively deceived. This particular skeptical problem will require that we produce some reasons or evidence that indicate that sensory experience is in fact reliable. To the skeptic, such evidence must directly rule out the possibility of massive deception. In other words, we must produce evidence that speaks against massive deception and in favor of the overall reliability of the senses. Such evidence then serves to justify our reliance on sensory experience. Further, without such justification, the skeptic will argue, that we cannot hope to ground knowledge claims in sensory experience.

It does not seem though that we embrace the skeptic's standards in everyday life. Rather, we seem to operate with *assumptions* of reliability for most of our major methods of belief formation (memory, testimony, sensory experience, etc). For example, while I have been able in many instances to verify the reliability of my memory, there have been a few instances in which my memory has been shown to me to be faulty. However, despite these few instances of error I have found via other methods of belief formation that my memory is generally quite reliable. However, I must admit that I have absolutely no confirmation that my memory has not been subjected to a skeptical scenario. If it were the case that a skeptical scenario involving massive and general deception were true, then my memory would be faulty and I would have no way to discern such massive error. In general skeptical scenarios, information fed to the agent in the scenario coheres in such a way that error is not detectable. What this shows is that I can't say with certainty that my memory is reliable, and, yet, I must admit that on many occasions I in fact continue relying on my memory knowing full well that is a fallible method of belief formation. There is a sense in which the skeptic is right, I do not know for sure that my memory is reliable, and, yet, on many occasions I use my memory to ground a knowledge claim. The skeptic is right that I tend to assume that my memory is generally reliable; I tend to assume that I am not a BIV. If we hope to do justice to the epistemic practices of the everyday believer then we must do justice to such assumptions of reliability.

First we must admit that the skeptic is right in thinking that we have no *evidence* to indicate that skeptical scenarios are false; that they do not in fact hold. Thus, we must admit that assumptions of reliability regarding sensory experience are just that-

assumptions. How is it that *assumptions* can play a pivotal part in justified knowledge claims? This is the fundamental challenge posed to the non-skeptic as I see it.

While I tend to agree with the skeptic that we cannot claim to know that we are not BIVs, I do not agree that our assumptions of reliability are not motivated at all. Thus, our defense of assumptions of reliability regarding sensory experience will involve a claim that admits that we are not epistemically justified in such assumptions, but that they are rational assumptions and as such can play a role in justified knowledge claims about the world. So even though we may not know that our senses are generally reliable, we are rational to assume so unless there is evidence to the contrary, and such assumptions are not inconsistent with claims to know.

Why might I think that assumptions of reliability need only be rational? The first reason is that assumptions of reliability are not directly justification conferring in our knowledge claims about the world. Particular instances of sensory experience are directly justification conferring. For example, my justification for believing that I have hands and my knowledge to that effect is my sensory experience of seeing and feeling my hands. Notice in many instances it is sufficient to ground a knowledge claim to the effect that such and such that I saw such and such. In the courtroom, no one questions that general veracity of sensory experience. Rather, it is sufficient to cite the relevant sensory experience.

Notice that what the above shows is that our assumptions of reliability regarding sensory experience operate as assumptions in our knowledge claims. We do not claim a stronger status than that. This is important. If we were masking assumptions as epistemically justified belief we would be in worse epistemic waters than we in fact seem

to be. We tend to treat assumptions of reliability as presuppositions in an argument, or as hinge propositions in a given context. The particular sensory experience is relevant to justifying a particular external world belief; the general assumption of reliability is not. Rather, the assumption is a presupposition to such a knowledge claim, but not that which directly justifies it.

So the view might be one that assumptions of reliability act as presuppositions that once in place allow the particular method of belief formation to act as *foundational* justification. Such a view has the benefit of explaining why it is only in philosophical contexts that such fundamental assumptions of reliability are called into question when defending a belief that is more widely defended by citing the particular experience involved in forming the belief. If particular instances of sensory experience enjoy a foundational status in the sense that they alone are sufficient to ground knowledge claims, then it is not clear what reason we have for thinking that the particular external world belief requires evidence in favor of the general assumption in order to be justified. In other words, if the particular sensory experience of seeing and feeling my hands is sufficient to ground my knowledge claim that I know I have hands, then it seems the skeptic's claim that my particular belief is not justified without further evidence in support of my acting assumptions is motivated at all. If the particular sensory experience is sufficient, then it seems that no further evidence is required in order to support the particular belief.

Instead of being in need of epistemic justification, assumptions of reliability only need be treated like any other presupposition in any critical thinking endeavor. That is, it need only be rational to have made said assumption. Since the assumption is not directly

justification conferring in the argument, it need only be a plausible or understandable assumption. More specifically we'll want to embrace the standard for all presuppositions. Presuppositions should be questioned only when we have reasons to think they are false. To criticize a presupposition on the grounds that it *might be* wrong are standards not in sync with the vast majority of our knowledge gaining and defending endeavors. This shows that the skeptic is stepping out of sync with the norm. This is relevant to us because the skeptic would like to think that she has taken the standards we ourselves claim to adhere to and shown that we in fact fail to meet those standards. She is right that we generally require justification to defend knowledge claims, but we do not require that all presuppositions required for the claim to be true must also be backed with evidence. There are long lists of necessary conditions that must hold for any given belief to be true. To require that everyday agents hold in their hand evidence that each and every necessary condition is true is to require even in theory that everyday agents are not capable of possessing true knowledge.

What I want to show here is not proof that the skeptic is wrong. Rather I want to more carefully show the nature of her insistence that knowledge is only consistent with complete certainty. What the skeptic demands is beyond the scope of the first person. Our epistemic practices and the methods of belief formation that we have been given cannot be verified from within our first person perspective, and, clearly, we cannot step beyond that perspective. For these reasons there is an inherent fallibility in our endeavors to seek truth and attain knowledge. The skeptic reminds me of someone who refuses to seek something so essential as food or love because she may not attain it, or even worse she may be fooled into thinking she's attained with no way of knowing for sure that she

has gotten hold of what she seeks. I want to defend the common sense and practical view that the pursuit of truth and knowledge is worth the inherent risk. Because the risk is inherent the start of the pursuit of truth involves ignoring error that might be only possible. We know from the get go that we do not have absolute certainty. But this admission is not tantamount to saying that we do not ever have reason to think we have in fact gotten hold of the truth. Our knowledge claims are not about claiming absolute certainty, they are direct claims to have the truth and they are founded on reasons to think we have attained the truth. Such reasons need not be infallible in order to be rational. In what follows I will attempt to defend the notion that assumptions of reliability are rational, they are not the epistemically irresponsible and psychologically convenient beliefs the skeptic makes them out to be. Further, I will try to articulate a sense in which such so called mere assumptions do not require us to give up our knowledge claims.

“Belief Formation in the Context of Everyday Life”

Chapter 3.1: Introduction

In this chapter we are going to apply the concept of epistemic perspectives towards a solution to external world skepticism. This solution is best characterized as a version of a relevant alternatives solution. Alternately, the solution could be characterized as a charge that the skeptic has committed the fallacy of accident. We will explore both characterizations in this chapter. Cartesian skepticism contends that an assumption about the reliability of the senses is not justified and is shown in the existence of skeptical hypotheses we cannot rule out. The argument in this chapter is that the de facto practice regarding sensory experience is that we are entitled to assume reliability and that such assumptions are epistemically sound and only dislodged given *evidence to the contrary*. Skeptical hypotheses are merely *possibilities* of error, and, thus not strong enough to dislodge our assumption of reliability. If our assumption of reliability for sensory experience is epistemically justified, then preference for our ordinary beliefs over their skeptical counter-parts is also justified.

It should be noted that the solution to be offered in this chapter will be derived from a discussion of the normative aspects of assumptions of reliability. As mentioned in earlier chapters, this discussion is aimed particularly at the worry that assumptions of reliability are not epistemically justified. Thus, we are taking aim at the skeptic’s

contention that the epistemically responsible, or the epistemically honest, position is that of the skeptic's. Those of us concerned with skepticism might worry that upon reflection, particularly in a philosophical context, the skeptic's argument seems to win the day. We might then wonder if the skeptic's position is the more intellectually honest position. These are the worries we are aiming at in this chapter.

The solution to be offered is not going to give a definitive solution in the traditional sense. The traditional approach demands that we give some evidence or reason that epistemically justifies a preference for our ordinary beliefs about the world over skeptical hypotheses. We are not seeking a solution in the form of refutation. Rather, I aim to show that the skeptic's position is not the only epistemically responsible position. The difference between the skeptic and non-skeptic seems largely to be a matter of differing epistemic values as much as anything else. The position of the non-skeptic is at least as defensible as that of the skeptic.

The above is the sort of solution being sought in this chapter. While such a solution might not rise to the level of a refutation of skepticism, it is also immune from the futile nature of attempting to *refute* skepticism. Many solutions have aimed to find *the* holy grail of philosophical arguments: a firm and sound refutation of the skeptic. Such refutation involves locating the much sought after evidence that shows ordinary beliefs are more likely to be true than their skeptical counter-parts. Yet, such arguments are notoriously flawed and often quite weak. Thus, rather than refute the skeptic, rather than attempt to quiet the skeptic, I aim to give those who are not skeptics a better sense of how and why their position is highly intuitive and epistemically defensible *even in light of uneliminated skeptical hypotheses*. It is likely that much of what I have to say will not

be compelling to the skeptic. However, I will draw out the epistemic implications of the skeptic's position and these consequences of skepticism should provide some relief from some of those pernicious first person worries that rise out of skepticism.

It might be useful to look at the discussion thus far. First, I have chosen to focus on one specific aspect of skepticism. Namely, our focus will be on the worry that knowledge claims are not justified in light of uneliminated skeptical alternatives. The charge behind this worry is that agents who claim knowledge are somehow epistemically irresponsible or epistemically irrational. Knowledge ascriptions have been characterized as "loose" ascriptions that are not literally true, or that knowledge in an everyday context represents "weak" knowledge.³⁶ It is this concept of skepticism I wish to poke a few holes in.

Next, it was decided in the introduction that skepticism has a distinctly *epistemic* aspect to it that is not addressed by stances on perception or theories in the philosophy of language. For the reasons given in those previous chapters I have chosen to approach skepticism within the traditional framework in epistemology. Additionally, the notion of epistemic perspectives articulated by Bonjour and filled out in the previous chapter serves to offer some concrete distinctions between the context of everyday life and the context of the philosophy classroom. While the distinction between everyday life/philosophy classroom and the first/third person perspectives is not precise it serves to illustrate some specific epistemic distinctions between belief formation in everyday life and the philosophy classroom.

³⁶ For an example of this sort of view see Peter Unger, *Philosophical Relativity*, selections from *Skepticism: A Contemporary Reader*, ed. by Keith DeRose et al (Oxford: Oxford University Press, 1999) 247.

For our purposes in this chapter we will want to keep in mind that everyday life usually involves first person concerns. Bonjour makes the distinction between first/third person perspectives by articulating what epistemic questions we tend to be concerned with. It also happens that when we are in an everyday context we tend to ask the same questions that define the first person and when in a philosophical context, we tend to be concerned with the questions that define the third person perspective. In everyday life, we are concerned primarily to *ascribe* knowledge and justification, and to assess such claims. A shift to the third person involves shifting to concerns such as the conditions for knowledge *in general* or justification *in general and from the perspective of an outsider looking in*, such a shift also often marks a shift into a distinctly philosophical perspective. The philosophical perspective is one in which the concrete parameters and the time limitations of everyday life have been put aside in order to pursue questions in a more thorough, systematic, and involved way.

In the next chapter, we will turn to the issue of epistemic closure and take a more detailed look at the issue of warranted assertability in relation to the view put forward in this chapter. It will turn out on the view to be articulated in this chapter that closure holds in a third person sense, but it sometimes fails in particular instances. The distinction is to be accounted for in the fact that closure in the third person is grounded in knowledge, but first person instances involve *knowledge ascriptions*. If the conditions for knowledge are distinct from those of knowledge ascriptions, it should not be surprising if closure holds as an abstracted third person form of reasoning but fails in many applications in the first person. Much of the next chapter will be devoted to

defending Dretske's claim that closure is a matter of the transmission of evidence and holds if the evidence for p is sufficient to claim knowledge of q, and fails if not.³⁷

Lastly, the focus of this chapter and much of the rest is to establish that we are in fact rational to rely on sensory experience. In this chapter, I will aim to develop the sense in which the pursuit of truth requires that we ignore some error possibilities. If we pay attention to any and all error possibilities, such practices often involve sacrificing the pursuit of truth. I will put forward the view that it is epistemically rational to assume reliability and treat as relevant only those error possibilities that are deemed sufficiently likely to warrant concern. An epistemic practice that allows prima facie assumptions of reliability to count as warranted epistemically seem most likely to be practices that maximize truth while avoiding too much error. The skeptic's position on the other hand, I will suggest, is an epistemic practice whose requirements on evidence are so strict and concern for error so deep that the skeptic sacrifices pursuing truth all because she is not certain of her faculties. Essentially I will suggest that truth is worth a leap of faith. The value of pursuing and attaining truth is so epistemically foundational that it warrants what is properly considered an epistemically rational leap of faith. Such a leap of faith is the beginning of pursuing truth and possibly attaining it. It is proper to deem the skeptic one who gives up possible truth out of her worry for possible error. The non-skeptic, alternately, allows for possible error because of the extent to which she values even possible truth.

In this chapter I will suggest that the constraints of everyday life and the realities of actual belief formation are such that we cannot entertain every possible error scenario

³⁷ See Fred Dretske, "The Case Against Closure." Contemporary Debates in Epistemology, ed. by Matthias Steup et al (Oxford: Blackwell Publishing, 2005) 13-25.

without sacrificing in a significant way our pursuit of truth. In the next chapter, I will try to show that relying on sensory experience is justified regardless of which possible world is the actual one. Specifically, I will aim to establish that since we do not know which possible world is actual, and we value truth, it is epistemically prudent or rational to rely on sensory experience even though it could be the case that we are in fact brains in vats.

Let us recall the version of skepticism discussed in the Introduction. Where H is a skeptical hypothesis and O is some ordinary belief about the world, Keith DeRose characterizes skepticism as follows:

- P1: I don't know that not-H.
- P2: If I don't know that not-H, then I don't know O.
- C3: Therefore, I don't know that O.³⁸

Thus, on this version of skepticism, our failure to count as knowing our ordinary beliefs about the world stems from (1) the fact that we don't know that not-H, and (2) the claim that our failure to count as knowing not-H implies that we fail to know that O. While some have attempted to derive a solution to skepticism by rejecting P1 and claiming that we do in fact know that not-H, we are going to focus on P2.³⁹

Notice, though, that attacking P2 via a relevant alternatives approach is a little awkward. Roughly, the relevant alternatives approach in epistemology is the view that only some competing alternatives need to be addressed by an agent claiming to know. To use a well worn example of Dretske's if one is at the zoo, looks in a pen and sees what appears to be a Zebra, then one might claim to know that there is a Zebra in the pen. Such a claim though is thought to be made against the backdrop of a limited set of alternatives. Thus, one would claim to know that the animal in the pen is a zebra as

³⁸ See DeRose, "Solving the Skeptical Problem," 2.

³⁹ See Gail Stine, "Relevant Alternatives, Skepticism, and Deductive Closure" for an example of such a solution.

opposed to a donkey or a horse, but not necessarily be claiming to know that the animal in the pen is not a mule cleverly disguised to look like a Zebra. Relevant Alternativists will argue that it is not part of the knowledge claim being made to claim to know *that*.

The kind of relevant alternatives approach found in Gail Stine's "*Skepticism, Relevant Alternatives, and Deductive Closure*" will refute skepticism via a direct denial of P1. The argument roughly goes that since "S knows that" is indexical the truth of the claim is relative to a particular set of contextually determined relevant alternatives. Once the set of relevant alternatives is considered, and, given closure for known entailment, it will follow that we do know that not-H. A more common strategy for attacking P2 is to deny epistemic closure under known entailment. My argument generally will go along the lines that skeptical alternatives are not relevant, but they are not relevant for epistemic reasons in the context of everyday life. Such a response amounts to ruling out skeptical alternatives on epistemic grounds.

Thus, it might seem that my argument involves a denial of P1. However, I do not wish to deny P1. Instead of denying P1 my argument involves the claim that we de facto treat skeptical alternatives as irrelevant in everyday life, and when evaluated in more depth this everyday practice is epistemically defensible, and, yet, the grounds of this defense are not sufficient to warrant claiming knowledge of not-H. Rather, the grounds for ruling out skeptical alternatives are sufficient to rule out skeptical alternatives, but not sufficient to ground a knowledge claim. Thus, on the view being advocated here what is required for ruling out an alternative is weaker than knowledge. The details of this issue will be explored some in the next chapter.

Thus, my argument, despite its being in the spirit of a relevant alternatives approach, will at heart involve some denial of closure. In the next chapter I will evaluate the relationship of my view to P1 and its implications for closure. It will turn out on my view, for reasons to be articulated later, that we know that O, and we know that O entails not-H, yet we do not know that not-H.

Let us begin by re-framing the skeptic's argument. The skeptic generates P2 by applying a certain principle or standard for evidence that we need to elucidate prior to evaluating. Let us call this principle the Rule of Exclusion. The rule governs justification in the context of forming beliefs when competing hypotheses are involved. We can characterize the Rule of Exclusion (E) as follows:

(E): For any phenomena P, if there are two competing hypotheses A and B and both are consistent with the evidence E regarding P, then an agent must have some epistemic reason for preferring either A or B over the other in order for her belief in either A or B to be epistemically justified.

Our intuitions regarding E seem to be as follows: if an agent prefers either A or B in a context similar to the one spelled out in E, then she prefers A or B for reasons that are non-epistemic and, therefore, her belief in A or B is arbitrary. Since beliefs that are preferred for non-epistemic reasons are epistemically arbitrary, they are not justified. Another way to think of this intuition is that we tend to regard beliefs that are not grounded in truth-conducive reasons or justifications as not *epistemically* justified. An agent that prefers A or B when they are both consistent with the evidence does so for something other than epistemic reasons. According to skepticism, it is counter-intuitive to ascribe knowledge to an agent when there is a competing hypothesis that she has yet to

rule out. If she cannot rule out H, then she does not *really* know that O. In fact, for all she knows, it could be that H.

Now we may spell out two scenarios, Scenario A and Scenario S, involving the application of E and re-characterize skepticism in the following way.

Scenario A: S is choosing between hypothesis A and B and her evidence E is consistent with both A and B. S chooses A for non-epistemic reasons.

We should conclude that S does not really know A, because S cannot rule out the possibility of B. In the context of skepticism the Rule of Exclusion may be applied in Scenario S in a way that is seemingly analogous to Scenario A:

Scenario S: Our ordinary beliefs O are in competition with a given skeptical hypothesis H, and agent X chooses O.

According to the Rule of exclusion, since X prefers O for non-epistemic reasons, she fails to count as knowing it. Given these two scenarios we may re-characterize skepticism.

We can think of this version of skepticism as skepticism via an argument from analogy.

We might characterize the argument from analogy as follows:

- S1: In most cases like Scenario A, the agent fails to know because she has violated the Rule of Exclusion.
- S2: The scenario of our ordinary beliefs competing with skeptical alternatives creates a scenario, Scenario S, that is relevantly similar to scenario A.
- S3: If S fails to count as knowing in Scenario A, then X fails to count as knowing in Scenario S.
- S4: S fails to count as knowing in A.
- S5: Therefore, X fails to count as knowing in S.
- S6: Since X fails to know in S, any agent relevantly similar to X fails to count as knowing.
- S7: Ordinary believers are relevantly similar to X.
- SC8: Thus, ordinary believers fail to count as knowing their ordinary external world beliefs that are grounded in their sensory experience.

The crux of this argument is S2, the premise that claims that Scenario A is relevantly similar to Scenario S. Given the above characterization of skepticism, we can attempt to

reformulate the skeptical problem as a question about what, if any, relevant differences exist between Scenario A and Scenario S. I want to explore the idea that there are some relevant differences between A and S such that applying the rule of exclusion in A is appropriate, but doing so in scenario S is a fallacy of accident.

3.2: Are there Relevant Disanalogies between Scenario A and Scenario S?

Even though we will look at a number of distinctions between A and S, the crux of my argument here will be that there are relevant distinctions between A and S that are particular to the first person perspective. The result of this argument will be a relevant alternatives solution in which we have some non-question begging reasons for dismissing skeptical hypotheses in many everyday contexts. In other words, I will argue in this chapter that there are norms in the first person perspective that (1) allow us to reject skeptical hypotheses and (2) have the result that the contemplation of skeptical hypotheses can actually result in the *violation of epistemic norms* that govern belief formation in the first person.

The primary thesis to be defended here is that error possibilities that are merely *logically possible* are appropriately irrelevant in the context of everyday belief formation. In other words, the merely logically possible is almost always irrelevant in everyday life. This conclusion follows in light of considerations involving the epistemic goal of truth and the concrete parameters of belief formation in everyday life. Thus, if this view is correct it should follow that treating the merely logically possible as relevant in the context of everyday life will result in the loss of true beliefs. One of the theses of this chapter is that such loss of truth for the sake of contemplating or avoiding merely

possible error is epistemically backwards. Such an approach is the result of prioritizing the avoidance of error over the pursuit of truth to such a degree as to sacrifice the pursuit of truth to the avoidance of error.

The argument in this chapter is going to rely on a number of considerations. First, the pertinent concrete parameters in everyday life are those involving time constraints driven by practical concerns that require beliefs be formed in a timely manner. The result of these concrete parameters is that the goals of pursuing truth and avoiding error conflict in the context of everyday life. Just as we have to sacrifice one moral duty to another in contexts in which two or more duties conflict, we often cannot fully pursue both of these two epistemic goals. Because these two epistemic goals conflict, we must sacrifice one to pursue the other, and we must balance the two concerns. This is the sense in which I am characterizing Cartesian Skepticism as a Moral Dilemma.

While there may be a number of things we can say about skeptical scenarios that make them distinct from other usual relevant alternatives, we are going to focus our discussion on two things in particular: (1) the concrete constraints in the first person that are in conflict with the contemplation of mere logical possibilities, and (2) the distinctly normative elements of justification in the first person perspective. Consideration of the first will generate the conclusion that it is for the sake of attaining truth that we focus on hypotheses that are *likely* and ignore alternatives that merely point out the *possibility* of error. In other words, it is for sound epistemic reasons that we operate with assumptions of reliability. Such assumptions of reliability are in place for most, if not all, of our fundamental mechanisms of belief formation. Thus, it will turn out that the argument

given here will be broader in scope and more general in its application than merely defending assumptions of reliability regarding *sensory experience*.

Considerations regarding (2), the considerations of the normative aspect of first person justifications, will offer a justifying account of everyday practices regarding sensory experience in the sense that agents are not blameworthy for their preferring their ordinary beliefs over skeptical alternatives. This subjective sense of justification will allow me to conclude, as have numerous others, that agents living in normal worlds and their skeptical counter-parts are not *blameworthy* for relying on their sensory experience. Thus, the arguments given here will offer a defense in both a subjective and objective sense of the epistemic practices of everyday knowers regarding sensory experience. This defense will ultimately need to be threefold. First, I will defend the everyday practice of treating skeptical alternatives as irrelevant. Second, since everyday agents are not epistemically neutral with regard to sensory experience, a justifying account *of this* practice is necessary. Third, we must defend the notion that agents are not blameworthy for relying on sensory experience. Defense of the first will be on the grounds that it is for good epistemic reasons that we are only concerned with alternatives that are probable and not merely possible in everyday life. Establishing the second will involve defending assumptions of reliability in general as conducive to our epistemic goals.

One way to think about the argument to follow is that it stands in contrast to Stroud's argument in "The Problem of the External World". Stroud argues that the standards the skeptic employs in a philosophical examination of our knowledge are no higher than those used in the course of everyday life.⁴⁰ The skeptic, according to Stroud, simply invokes the standards that are at use in everyday life to get her skeptical

conclusions. The question about the general reliability of our senses stems from Descartes' desire to assess all of his seeming knowledge, and the best route for this is to assess those *general* sources or methods by which one forms beliefs. If *those* sources are found wanting, then our beliefs themselves cannot amount to knowledge. Stroud gives an analogy: "If I found good reason to doubt the reliability of the suspect's alibi, for example, and that was all I had to go on in my belief that he was in Cleveland, then what I earlier took to be my knowledge that he was in Cleveland would have been found wanting or called into question. Its source or basis would have been undermined."⁴¹ In a similar way, if the sources of our beliefs are found wanting, then our "knowledge" has been undermined.

I accept the above line of reasoning in the argument to follow, and it seems correct to think that the crux of skepticism lies in our reliance upon sensory experience; this primary source of information about the world is under review and found wanting in the skeptic's mind. There are two challenges posed to this line of reasoning in the argument to follow. First, we are going to reject the notion that skeptical hypotheses constitute a "good reason" for questioning sensory experience. Let's grant Stroud, that *if* we have a "good reason" for doubting the source of a belief, then the belief itself is undermined. However, I will argue that it is not as clear as the skeptic would like that skeptical hypotheses constitute "good reasons" for doubt that in fact undermine the source in general. The heart of this argument will center on the claim that skeptical hypotheses are merely logically possible, and possible error is not the same as likely error. Thus, my argument will go that the better analogy lies not in comparing having a

⁴¹ Barry Stroud, "The Problem of the External World," *Epistemology*, ed. by Ernest Sosa et al (Malden: Blackwell Publishing, 2008) 9.

good reason to doubt an alibi, but in recognizing the logical possibility despite the quality of one's evidence that the alibi is not good. Thus, the appropriate analogy is between recognizing the possibility of error in relying on an alibi and the possibility of error in relying on sensory experience.

Second, the argument to follow in later chapters will reject the notion that we must have reasons or evidence to indicate the reliability of a method of belief formation in order to be justified in relying on it. Thus, I will be arguing that skeptical arguments like the one characterized above are based on two less than clear premises: the first is the claim that skeptical hypotheses constitute good epistemic grounds for doubting sensory experience, and (2) that we must have some positive reason to rely on a method of belief formation in order to be justified in doing so. Premise (1) is clearly stated in Stroud's argument. Premise (2) is the premise that would be needed to get a skeptical line of reasoning going if we reject (1). The remaining argument in this chapter will argue against (1), and the arguments in later chapters will aim more specifically at (2) - the assumption that we need some positive evidence in order to be justified in relying on a method of belief formation.

3.2.1: Skeptical Alternatives are irrelevant because the merely possible is almost always irrelevant in the context of everyday life:

Notice that in most contexts in which there are competing hypotheses, we usually have *some reason* to think the alternatives are either plausible or true. Additionally, in most scenarios like Scenario A we rely heavily on our background beliefs to help us determine which of the hypotheses seem most likely. Skeptical hypotheses and our ordinary beliefs are not equally supported by the evidence – they are so equally supported

only if all other beliefs are taken out of account. There is some precedent in the literature for thinking that what it takes to get into a philosophical context is relatively simple. If someone raises a skeptical hypothesis, then we are in a skeptical context. While it does seem that there are contexts that are distinctly skeptical, contexts in which a skeptical alternative is under consideration, what it takes to be in such a context is arguably more stringent than often recognized.

Lastly, that skepticism does not fit well into the context of everyday belief formation- that it is so difficult to *practice* this philosophy indicates that it is a concern appropriate only for the philosophy classroom. It is possible to take skepticism seriously only when we have removed ourselves from the context of making decisions and judgments about what we immediately have in front of us. Another way to make this claim is that even for philosophers and skeptics it only seems appropriate to contemplate skepticism in a philosophical context. While it might be in the tradition to assume this is so because of psychological reasons beyond our control, I am inclined to think of this response as an ad hoc recovery attempt on the part of the skeptic. It seems our attitude, rather, is that it is not epistemically *appropriate* to take skepticism too seriously in the context of everyday life. This everyday attitude warrants further evaluation, because it lies at the heart of the skeptic's attack on our claims to know. If we are to truly defend everyday knowers from the implicit charge of irrationality, we have to defend the practices of everyday knowers in a way that is consistent with their actual epistemic practices. It is for this reason that we cannot jump to the arguments contained in the final chapter and ignore the arguments given here. Fancy philosophical reasoning that everyday knowers have never appealed to does little to defend *their* claims to know.

The view that we rely on sensory experience because we are hardwired to do so does not capture a prevalent attitude towards skeptical alternatives. The prevalent attitude is that skeptical alternatives are not relevant and that everyday agents are not under any epistemic obligation to respond to these possibilities. Thus, the “psychological compulsion” explanation of our everyday attitude to skepticism does not fit the data to be explained. The explanation needs to involve some account of why we tend to think it is *epistemically inappropriate* to take skeptical hypotheses seriously in the context of everyday life. One of the aims of this paper is to explore and defend *this* attitude towards skeptical hypotheses.

To see the former, let us perform a thought experiment- call it thought experiment G. Let’s imagine that a government is attempting to assess evidence regarding the activities of a rival nation in order to assess the threat level to national security. There is a body of evidence that has been amassed regarding the activities of the rival nation, and intelligence agents are developing hypotheses and scenarios that explain that body of evidence. So each agent develops a hypothesis about the activities and motives of their rival that explains the data that has been accumulated. Now let us imagine that one of the agents develops a hypothesis that is merely a logical possibility. The agent has no reason to think her hypothesis is likely, the hypothesis does not fit any of her background beliefs about the rival, and yet she presents her scenario as a contending hypothesis. When asked *if she herself* believes, or prefers (if belief is too strong), this scenario she states that she does not.

To make our example more concrete we might imagine the agent’s hypothesis is that aliens have created and faked the amassed data and disconcerting evidence with the

goal of creating strife on Earth. Of course, she is then asked to defend her presentation of this scenario and her defense is that her hypothesis is consistent with the evidence, and it is in fact *logically possible*. The only appropriate response is not only to dismiss her scenario as having a very small degree of likelihood, it would be appropriate to be angered about her presentation. Valuable time and money and effort have been wasted. These agents are in a context in which decisions must be made and made within a certain framework. They are in a context in which *mere logical possibilities* are not only irrelevant but inappropriate and destructive to the epistemic constraints and goals in the context. They are concerned with what is *likely*, not what is possible, and, thus, in order for the agents to fulfill their epistemic duties in this context they must develop hypotheses that they themselves prefer and *believe to be likely*. They must have some reason for thinking their alternative is the most likely explanation.

The agents in the above example are focused primarily on arriving at truth, and the standards for adequate evidence are relatively high given the serious nature of the inquiry. It is the concrete time limitations that prevent them from entertaining every possible explanation, not their pragmatic or non-epistemic goals whatever those might be. It is because of such time constraints and the goal of arriving at a well-founded hypothesis that they ignore the merely logically possible. They cannot *both* entertain every option and come to a *likely* conclusion within the given time frame. So, it is for the sake of pursuing truth that some possibilities are ignored. One way to formulate the thesis of this argument is that the skeptic is like the agent in this example. She presents

mere possibilities in contexts that require contending alternatives to be on stronger epistemic grounds than merely *logically possible*.⁴²

The above is the sense in which the skeptic commits the fallacy of accident. She has applied the Rule of Exclusion to a set of alternatives for which it was never meant to apply. In general the rule of exclusion requires agents to have evidence or reasons to rule out contending alternatives in order to be justified in preferring one over the other. And, in general, one is required to have evidence that *speaks against* the contending alternatives. However, it is also true that the Rule of Exclusion is meant to apply for well-founded alternatives. The reason for this is clear. First, we tend to think that agents who ignore well-founded competitors are doing so for non-epistemic reasons; this motivates the worry that such agents are indeed being irrational and irresponsible in their belief. However, it also seems to be true that we do not apply the Rule of Exclusion to each and every alternative hypothesis; we do quickly reject alternatives that we deem far-fetched or insufficiently grounded. This latter phenomenon characterizes skepticism better than the scenario in which an agent presents a well-founded alternative only to be dismissed without sufficient reason.

Notice that I am not making the contentious claim that skeptical hypotheses are *unlikely*. Nor am I committed to the contentious claim some relevant alternatives solutions have used to ground their arguments. The contentious claim is that skeptical alternatives are not relevant because they are *unlikely*, and they are unlikely because they

⁴² However, to be fair, this is too hard on the skeptic. My own sense is that skeptics do not actually present skeptical hypotheses in everyday life – they are presented in philosophical contexts. The mistake is made when we philosophers go along with the skeptic in thinking that raising relevant alternatives is as simple as mentioning them - with the result that skeptical hypotheses can easily become genuine competitors to our everyday beliefs. The result of this hasty move is that we philosophers develop the worry that our everyday beliefs have been undermined by skepticism. The argument being pursued here is that such is not the case.

occur in *remote* possible worlds. Such a view is seen in Keith DeRose's claim that he has hands is sensitive betrays his conviction that skeptical worlds are remote possible worlds, thereby allowing him the conclusion that his claim to have hands is sensitive with regard to all relevant possible worlds.⁴³ If we include our background beliefs in the assessment of skeptical alternatives, then it seems we are justified in treating them as *unlikely* contenders. They are unlikely because our background beliefs tell us such a hypothesis is far-fetched. However, such a move involves allowing biased beliefs to factor into rejecting skeptical alternatives. The skeptic has a legitimate criticism of begging the question if we were to make such a move.

Rather, my argument is grounded in the claim that skeptical hypotheses are logically neutral – they are neither likely nor unlikely. There is no reason or evidence to believe that skeptical alternatives, such as the BIV hypothesis, do in fact hold, nor is there any reason or evidence to believe that it does not hold. Skeptical hypotheses, therefore, are merely logically possible, and, as such, they are ignored in the context of everyday life. My argument is that such an attitude towards the logically possible is both warranted epistemically and required pragmatically.

In order to substantiate this position further, a little should be said about why the skeptic's sense of what is relevant is not appropriate. What we are concerned about here is the possibility that the skeptic can claim that *we should* be concerned about her skeptical scenarios, and that we should be so concerned, *regardless* of what she takes to be the *merely practical* constraints of the context. In order to address this we need to establish two things: (1) what the skeptic takes to be *merely practical constraints* are epistemic as well, such that the *norms stemming from the practical constraints are*

⁴³ Keith DeRose, "Solving the Skeptical Problem," 18.

epistemic rather than pragmatic, and (2) that knowledge is possible within the first-person context of everyday life. Arguing for (1) and (2) should allow us to block the above possible move on the part of the skeptic. First, the skeptic might contend as many have been inclined to that our reasons for ignoring skeptical alternatives are grounded in the non-epistemic, and, therefore, cannot serve as epistemic justification. To argue for (1) we need to show as I have argued thus far that our reasons for ignoring skeptical alternatives are epistemic. However, the argument as it stands offers some kind of subjective sense of justification. The latter part of this chapter and the bulk of a later chapter will focus on substantiating the claim that an assumption of reliability regarding sensory experience results in the attainment of more truth.

Secondly, the skeptic may want to argue that the practical constraints of the first person context prevent the attainment of *true* knowledge. The skeptic might be inclined to say in response to the argument above that since we cannot fully avoid error, then we cannot pursue or attain true knowledge in the context of everyday life. I have presented the skeptic as requiring that we must rule out all possibilities of error in order to count as knowing. However, I have argued that to adopt such a standard will amount to avoiding error at the cost of pursuing truth with the conclusion that the skeptic's standards are obviously too high and too stringent. If we must choose between avoiding all possible error or the pursuit of truth, we should choose the pursuit of truth. However, there is nothing to prevent the skeptic from arguing the other way. She could very well argue that if we cannot avoid error to the degree she thinks necessary in order to count as knowing, then she may be inclined to conclude that true knowledge is not possible in the context of everyday life.

The bulk of this chapter and the next are spent offering an epistemic defense of our assumption that sensory experience is reliable. Therefore, I want to address this second potential response on the part of the skeptic before returning to the rest of the argument at hand. Admittedly, this second potential move on the part of the skeptic is more difficult to address. In terms of the semantics of the verb, a skeptic might be inclined to argue that “to know” is an example of a term that might apply loosely in certain contexts but whose truth-conditions are never literally met. If this is the crux of skepticism as it may well be, then we might find ourselves at a loss for ways to respond on the approach I have taken here. If the skeptic simply stipulates that knowledge requires certainty, this will imply that an agent must be able to rule out all possibilities of error for definitive reasons, which will include the ability to rule out skeptical hypotheses. Clearly, on such a definition of knowledge, knowledge in everyday life will be fairly unlikely. On such a definition, knowledge, if possible at all, will be the result of much thought and work outside of the constraints of everyday life. So that it will be mathematicians and philosophers (on occasion) that will count as possessing knowledge. The disagreement we might have with the skeptic here will stem from fundamentally different intuitions on what knowledge is.

However, there are a few things we *can* say in response to such a move. It is possible to make arguments to the effect that the skeptic’s intuitions are out of line with how most use the term “knowledge”. There is a large body of evidence that indicates that as a matter of fact we accept knowledge as consistent with the possibility of error. Nonetheless, there have been philosophers who have notoriously rejected this standard of knowledge in favor of more stringent standards. Having said this, I still have a strong

sense that if philosophers or skeptics stipulate that knowledge implies certainty, which itself implies having strong grounds for ruling out skeptical hypotheses (after they have been introduced), then we have the right to question such a stipulation. Such an epistemology dismisses by definitional fiat much of what is of interest in epistemology and reduces a concept that appears to have fundamental importance in our real lives to a classroom novelty. To my mind, there is enough that is unappealing in this approach to justify trying to avoid it. Such a response will require that we not allow or not get on board with a skeptic who insists that knowledge implies certainty.⁴⁴

Notice that such a skeptic will be committed to a much broader skepticism than external world skepticism. If the argument that drives the skeptical conclusion regarding sensory experience is that sensory experience is not the kind of evidence that allows one to rule out all error possibilities with evidence that speaks against them, then such skeptical reasoning will apply to most of our methods of belief formation if not all of them. Descartes reasoned along these lines for much of the first Meditation, and if we follow this line of reasoning, we will end with the same kind of global skepticism that Descartes found himself in, and we are likely to struggle just as much to get out of such global skepticism. It is just a fact that most if not all of our methods of belief formation are fallible; even a priori reasoning is fallible. Some have been inclined to conclude on the heels of such reasoning that knowledge is a term whose truth conditions are never literally met. However, such a move seems justifiable only on the heel of failure to account for knowledge as consistent with fallibility. If such an account can be given,

⁴⁴ Perhaps we want to allow for strong knowledge and weak knowledge. If the conclusion of skeptical arguments is that we often do not know in a strong sense, this does not seem problematic to me. On such a view, what we would be concerned about is any skeptical argument to the effect that we fail to possess even weak knowledge of the external world.

then there is no motivation to characterize the semantics of “to know” in a way in which the truth conditions are never literally met.

3.2.2: Assumptions of reliability are *epistemically justified*:

If we allow skeptical alternatives as relevant, the skeptic believes we are at a stale-mate. The usual evidence is consistent with both our ordinary beliefs and skeptical alternatives. Thus, certain breeds of skepticism have encouraged an epistemically neutral stance regarding sensory experience. This stance regarding sensory experience is purportedly the only epistemically responsible or rational position in light of the seeming stale-mate. If this skeptical reasoning is sound, it will turn out that everyday agents have poor epistemic practices because they generally *assume* the reliability of sensory experience.

In this section we must defend the idea that skeptical alternatives can be ruled out simply on the grounds that mere possibilities of error are not sufficient to overwhelm a warrant to assume reliability. First, let’s note that the skeptic wants to have her cake and eat it too. She wants *to raise* a skeptical hypothesis without any evidence of its likelihood, and, yet, wants the stricter standards of evidence *for ruling out* her hypothesis. The stricter standard would be to require *evidence that speaks against* the hypothesis as the only sufficient grounds for ruling it out. This criterion would amount to the requirement that an agent must *know that not-H*, where H is some skeptical hypothesis, in order to rule it out. However, we can ask why it is that the skeptic is entitled to raise alternatives on such weak grounds and expect such strong grounds for ruling out her

alternative. It is not as though raising alternatives in the context of everyday life is without its consequences and implications for the pursuit of truth. Thought experiment G demonstrated that the contemplation of an alternative uses resources with the implication that doing such can hamper the epistemic goal of pursuing truth. Thus, we might conclude that it not only possible to be irresponsible in ruling out alternatives, it is possible to be epistemically irresponsible in raising them as well.

Let us look in detail at what might follow if we were to practice a skeptical stance with regard to sensory experience. If we are not entitled to assume reliability, then it will follow that we either violate our epistemic responsibilities or take an epistemically neutral stance with regard to sensory experience. If there are problems applying or practicing a theory or philosophical position, it is good reason to think that there are problems in the theory itself. What might follow if we are not entitled to assume the reliability of the senses? First, note that if we were to do such it would amount to sacrificing the pursuit of truth regarding our external world for the sake of avoiding all risk of error. While we are concerned with error for good reason, we are not so concerned as to give up the pursuit of truth. This is so particularly in circumstances like those of skepticism in which there is no *positive reason* to think we are in error.

Let us imagine what we might say if we were to grant the skeptic her worries. What concrete proposal follows in light of this concession? Shall we stop relying on sensory experience? Such a proposal seems difficult, if not impossible to follow, and, frankly, like quite bad advice. Perhaps the skeptic thinks a more subtle shift is in order.

Perhaps we should simply stop *claiming* knowledge on the basis of sensory experience.⁴⁵

I find this approach somewhat inconsistent if not superficial as argued in the previous chapter. Relying on sensory experience should not be a mere practical matter of doing what is psychologically compulsive. The reliance itself embodies the *epistemic* stance that it is a method by which one comes to truth.

It is the *reliance* upon sensory experience itself that embodies this epistemic stance- not the claim to knowledge. The claim to knowledge merely comes on the heels of the assumption of reliability. Relying on sensory experience to determine our beliefs about our environment is no small practical matter. It is a matter of central importance in our *epistemic* lives. What beliefs are more central and fundamental than those we form about our immediate environment? We rely on sensory experience with the belief that we are attaining truth. The *claims* to knowledge are a secondary and less important reflection of this more fundamental epistemic stance of endorsement. Thus, if relying on sensory experience is a violation of our epistemic duties with the conclusion that agents who do so are irrational, then refraining from claiming knowledge does very little if anything at all to save agents from such epistemic failure.

Returning now to the first possible move, the skeptic might contend that our dismissal of skeptical hypotheses is being made on pragmatic rather than epistemic grounds. If we recall the first possible move the skeptic might make it to insist that it is for pragmatic reasons that we refuse to count skeptical alternatives as relevant in the context of everyday life. So let us review the reasons for thinking that our dismissal of skeptical alternatives is grounded in the epistemic.

⁴⁵ Such an approach would be in the spirit of Pyrrhonian skepticism. Since we do not know whether or not we know, we had best suspend judgment on the issue. Thus, neither claiming nor denying knowledge is appropriate on this line of reasoning.

Imagine the effects for attaining truths if everyday agents actually took seriously each and every logical possibility. The effect of such a practice would be that everyday knowers would start to look a whole lot like philosophers. Knowers would begin to feel as though they need to rule out and address every possibility prior to endorsing something. If we imagine this practice in real life, the consequences for the attainment of truth are profound. If everyday knowers did not have the entitlement to dismiss mere logical possibilities, it would come at the cost of the attainment of many possible truths. The epistemic goal of avoiding error cannot come at such a huge cost to the goal of attaining truth.

The result of this deeply embedded hierarchy of epistemic values is that it is for *epistemic* reasons that we risk some possibility of error – it is for the sake of possibly attaining truth that we take a leap of faith. It is for epistemic reasons that we operate with an assumption of reliability for most of our fundamental methods of belief formation. A *possibility* of error is not sufficient to dislodge this assumption of reliability. Rather, we will want to endorse the looser requirement that we need only question our method of belief formation in light of evidence that said method is in fact unreliable or in light of evidence the existence of which an agent should be aware of. Of course, *the skeptic* has not met *this* standard.

If this were not the case, if we did not form beliefs in the context of everyday life and within the very concrete parameters of our environment, we might pursue truth in a way that looks very different. In fact, this is precisely what philosophers do. It is *philosophers* who broaden the scope of what is relevant in order to take a fuller view of what is possible and what might be. But we should not assume that practices appropriate

for the philosophy classroom are epistemically appropriate in the context of everyday life.⁴⁶

Notice that we need not be committed to denying evidentialism.⁴⁷ Evidentialism is the view that if two agents possess the same evidence with regard to belief *p*, then they stand in a similar epistemic situation with regard to *p*. Denying evidentialism amounts to claiming that two agents can possess the same evidence for *p*, and, yet, be on different epistemic grounds with regard to *p*. Such divergence would need to be accounted for in terms of the pragmatic differences between the agents' contexts. If we allow pragmatic differences to affect epistemic standing, then we have included a distinctly pragmatic component into the concept of epistemic justification.

While my view might easily seem to be arguing against evidentialism, I do not think it necessary that it should. I am inclined to factor the concrete parameters of the everyday context in somewhere other than in terms of epistemic justification. We need not claim that the evidence the everyday knower possesses is not sufficient in the philosophy classroom, such that the everyday knower is justified in her belief that she has hands but the philosopher is not so justified. Rather, it seems perfectly plausible to claim that both are equally epistemically justified, but the philosopher has spent more time analyzing and understanding her justification. Such reflection might have other benefits rather than resulting in a *fundamentally* stronger epistemic stance. Perhaps the result of

⁴⁶ Since I have already engaged in a little philosophy-bashing here, we might as well go all the way. I suspect that it is a very good test of a propensity for philosophy to bring up a skeptical alternative in class and see how students react. Those students who are concerned and paralyzed by the skeptical hypothesis show a strong propensity to philosophize. They are willing to expand the boundaries of what is relevant. On the other hand, those students who find it silly to spend time thinking about such wild hypotheses show a good dose of common sense, but would probably make for very bad philosophers.

⁴⁷ See Feldman and Conee, "Evidentialism," 310-322.

philosophical reflection is added comfort and security resulting in the knowledge that a decision made quickly was ultimately the right decision after all.

Alternately, we might endorse Foley's distinction between epistemically rational belief and responsible belief. Where epistemically rational belief is tied to epistemic justification and responsible belief involves epistemic justification as well as all other pragmatic issues that factor into belief. Responsible belief, according to Foley, is belief that one's procedures with respect to *p* have been acceptable in light of limitations on time, and in light of *all* of one's goals. On such a view then we might be inclined to say that what amounts to responsible belief for the everyday knower, who faces time constraints the philosopher does not, does not amount to responsible belief for the philosopher. However, both the philosopher and everyday knower are justified in their belief that they have hands grounded in the sensory experience of seeing and feeling hands. Thus we would characterize both as having epistemically rational belief.

In *Knowledge and Practical Interest* Jason Stanley offers reasons to treat a contextualist position that treats knowledge as indexical with some prima facie suspicion. He argues convincingly that "knowledge" does not behave in the same way that uncontested indexicals do.⁴⁸ For example, with classic indexicals it's not problematic to shift standards within a given sentence or a brief exchange. However, it does not seem so easy to shift standards for "knowledge", and if such standards did shift in the way that classical indexicals do, then it would be problematic for those philosophers who wish to use such shifts to diagnose skepticism and offer a solution. If it is not inappropriate to shift the standards in a given exchange, then the anti-skeptical conclusion of a

⁴⁸ Jason Stanley, "Knowledge and Practical Interest," selections in *Epistemology: An Anthology*, ed. by Ernest Sosa, et al (Malden: Blackwell Publishing, 2008), 727.

contextualist position is lost. Thus, if “knowledge” does behave like other indexicals, then a contextualist view based on such a position would have to be more concessive to the skeptic than previously thought.

Largely for the above reasons, I do not wish to endorse a contextualist view of knowledge. It seems that such purported shifts in the standards for knowledge would likely be accounted for in shifts of what counts as “adequate evidence.” It should be noted on the view I am endorsing here, the standards for knowledge possession remain the same. Perhaps the standards for knowledge ascriptions shift slightly when we shift from an everyday context into a philosophical one, or shift slightly when the pragmatic factors in the context shift (i.e. if we change into a context in which there is a lot riding on the belief in question). However, one of the aims of this chapter is to show that even in the context of knowledge ascriptions, the standards do not shift easily. The skeptic cannot shift standards merely by *mentioning* her skeptical hypothesis. As argued, it seems to me that in most such contexts, people quickly dismiss that which is merely possible.

Notice the brand of fallibilism about knowledge that is on the table here. As Lewis notes in *Elusive Knowledge* it just sounds weird to state that “He knows, yet he has not eliminated all possibilities of error”⁴⁹ One of the appealing aspects of the account being formulated is that while we are endorsing fallibilism in a sense, we are not endorsing this very counter-intuitive version of it. At this point we want to characterize our everyday knower as having rejected all other contenders for epistemic reasons. The agent is in an epistemic state in which she does not believe there are possibilities of error

⁴⁹ David Lewis, “Elusive Knowledge.” *Skepticism: A Contemporary Reader*, ed. by DeRose et al (Oxford: Oxford University Press, 1999), 221.

that she has not ruled out. However, we will grant that her grounds for ruling out some hypotheses are not definitive or certain. This is the element of fallibilism in the account. The everyday knower as we have imagined her is an agent that has looked at all serious alternatives and ruled them out (or, if she has not done this in practice – she could if called upon to do so).

However, unlike other many other contextualist and relevant alternatives solutions, our agent could rule out skeptical hypotheses if called to do so. Her defense of her preference for her ordinary belief is grounded in her reliance upon sensory experience. The response that a mere possibility of error is not sufficient to justify a change in her stance regarding sensory experience is an epistemically sound response in light of the fact that skeptical alternatives embody error that is merely possible. The skeptic is articulating an alternative approach, that the epistemically responsible stance is to not endorse any method that has not been proven reliable, but we are arguing that this alternative is not sound epistemic practice.

The context of the first person is one in which the context dictates epistemic norms that emphasize speed and efficiency in order to attain a sufficient amount of truths. Thereby, the context of the first person is one in which hypotheses that are seen as probable are those that are relevant and cannot be ignored. Therefore, it is the goal of attaining truth that underpins assumptions of reliability. We cannot both pursue truth and entertain every error possibility; we must make a choice and sacrifice the concern for error that is merely possible to the goal of attaining truth. Of course it is not as though we abandon the concern to avoid error all together. Error that seems likely cannot properly be ignored. Thus, error for which we have reason to believe has occurred or is

likely to occur cannot properly be ignored; doing so results in violating the duties regarding avoiding error.

3.2.3: Agents are not blameworthy for relying on sensory experience:

Agents are not blameworthy for a number of reasons, many of which involving notions of moral blameworthiness in general. If we can show that even BIVs are not blameworthy for relying on sensory experience, then it will follow that normal agents are not either. This will go the other way as well, if BIVs are blameworthy, then so are normal agents. Blameworthiness cannot turn on the luck of which world is the actual one in light of the fact that it is impossible for agents to tell which possible world is actual.

Let us turn to the concern about whether or not agents are not epistemically blameworthy for relying on sensory experience. Note that in the first person justification is largely normative in nature and our epistemic practices seem to be largely guided by the intrinsic value of truth. Agents are not blameworthy for assuming the reliability of sensory experience because such an epistemic practice is highly conducive to the attainment of truths. Clearly, we cannot sincerely accept in practice that since our senses have not been *proven* to be reliable we should not rely on them. The value of truth, like all values, must be balanced against other competing values—such is the case in much of our moral lives and such is the case in our epistemic endeavors as well.

The skeptic is asking to put her alternative on the table in the first person perspective to have us address it using first person justification (those things we can access via reflection alone). Yet, we have shown that skeptical alternatives are really not like usual alternatives. It is impossible to assess a skeptical hypothesis in the first person

because there are no indicators for or against the hypothesis. The differences are very relevant and, therefore, there are important disanalogies between Scenario A and Scenario S. Both Scenario S and Scenario A involve norms about belief formation, in the sense that certain ways of forming beliefs are seen as responsible and others are not. The intuition that supports the Rule of Exclusion is an intuition about when belief is responsible and when it is not.

Another way to frame the argument thus far is that there are practical matters that inform our epistemic attitudes and practices. While I do not wish to go as far as some, the argument given here does not even require a contextualist account of knowledge, it does seem that practical matters have informed the general attitude we have taken towards sensory experience. It seems that our endorsement of the reliability of sensory experience only needs to be *rational* and epistemically justified *in light of* the practical circumstances in which we live. There is little that compels one to accept or endorse anything stronger than this. I have argued that endorsing sensory experience, and, our assumption of its reliability is epistemically sound.

3.3: Other Relevant Alternatives Views:

Other relevant alternatives views have struggled to articulate when and how skeptical alternatives become relevant. The consensus seems to be that if the agents of the context accept the shift in standards that allowing skeptical alternatives involves, then the agents will shift into a skeptical context.⁵⁰ So the standard view seems to be that

⁵⁰ Such a view is seen in Duncan Pritchard's discussion of knowledge ascriptions in Epistemic Luck, David Lewis's "Elusive Knowledge", Stewart Cohen's "Contextualist Solutions to Epistemological Problems," Gail Stine's "Skepticism, Relevant Alternatives, and Deductive Closure", as well as Dretske's "Epistemic

someone in the context raises skeptical hypotheses, and the others tend to engage the hypothesis or not. Some views make it look as though the mere mentioning of an alternative forces the context to become a skeptical one. This follows from the fact that such views tend to endorse a contextualist view of knowledge and on such a view, the standards for knowledge shift. It has turned out on some views that skeptical alternatives are not relevant because of the happy accident that no one has raised such an alternative in the context. Many contextualist views indicate that raising an alternative is as easy as mentioning the alternative. Duncan Pritchard gives a contextualist reading of knowledge ascriptions in his *Epistemic Luck* and in this such views are quite vulnerable to skepticism and largely concessive to it. The skeptic need only mention her skeptical alternatives in order to refute our claims to know. Though it may not follow that we fail to know in ordinary contexts, such is the result of too much that is accident and luck.

However, if we find epistemic grounds for treating skeptical alternatives as irrelevant, then perhaps we have moved away from the spirit of the relevant alternatives view in general. Such might be said of the view I have offered here. Rather than being a relevant alternatives view, it might be said that I am simply offering reasons for rejecting skeptical alternatives, thus I need not rely on a notion of relevant alternatives. I have characterized my view as a relevant alternatives view because it is my sense that our de facto position in everyday life is that skeptical alternatives are not relevant. They are not part of what we consider in everyday life. This is the sense in which my view is a relevant alternatives approach. Additionally, it has not been my goal to claim that we know that not $\neg H$, where H is any given skeptical hypothesis. Instead, my claim is that H

Operators” to name a few. Each of these views discusses the skeptical context as though it is one we get into by having a skeptic raise a skeptical alternative via discussing the alternative.

is not relevant. Where a claim to know that not-H requires some evidence/reasons in favor of that specific hypothesis, a claim that not-H is epistemically irrelevant or inappropriate in the context is a different and epistemically weaker claim.

If we were to develop a full account of conditions in which alternatives may properly be ignored, one of those conditions would be that the merely logically possible is properly ignored on the straightforward grounds that the logically possible is not sufficiently grounded to justify the time it takes to consider it seriously. This will follow from the fact that the time it takes to consider an alternative seriously is time taken from the pursuit of truth and the contemplation of hypotheses that are on stronger grounds. There are realities to the formation of belief and the attainment of knowledge in the real world that we philosophers have been too inclined to ignore or dismiss. However, it is in these real-world considerations that a solution to skepticism lies.

3.4: Assumptions of Reliability about Sensory Experience are Truth-Conducive:

That the above practice is grounded in the epistemic rather than the pragmatic results from the following: it is for the sake of truth that we ignore error scenarios that are merely possible. The argument in this section is an outline of the argument in the next chapter. The argument in the previous sections of this chapter was meant to establish that from our first person perspective, in a subjective sense, we ignore possibilities of error for the sake of pursuing truth. My contention in this section to be filled out in the next chapter is that doing so results in the attainment of more truths. This will imply that assumptions of reliability are objectively truth conducive. Admittedly, we are forced to alter our analysis of truth-conducive from the traditional use

of this concept. The details of this shift and the reasons for it are defended in the next chapter, but the gist of that argument is outlined below.

Let's imagine that agents in all possible worlds operate with assumptions of reliability regarding sensory experience. If agents in both normal and skeptical worlds rely on sensory experience, then agents in normal worlds would attain more truths by so doing, while agents in skeptical worlds would attain more false beliefs.

On the other hand, if all agents act in accordance with skepticism and endorse only those methods that have been *proven* reliable, then agents in both skeptical and non-skeptical worlds will remain epistemically neutral with regard to sensory experience.

The result of *this* practice would be that agents in skeptical worlds would have significantly less false beliefs, but, notice that their restraint would not result in more true beliefs either since they have not been provided a route to such empirical truths.

Empirical truths are not possible truths for skeptical agents; such is the result of the mechanism of deception in skeptical scenarios. Note also that agents in normal worlds would be *missing out* on their possible true beliefs because they have followed the skeptic's strict standards against assumptions of reliability. Thus, while there are less false beliefs in this scenario, this practice results in the loss of truth as well. Indeed, in this scenario, *no one* is attaining truths about themselves and their environment.

It will follow that relying on sensory experience results in more true beliefs. Granted, it may not result in more true beliefs *over* false ones, but it results in a greater amount of truth when compared with the skeptic's alternative approach regarding sensory experience. Additionally, notice that assuming the reliability of the senses generates truth in a way that refraining from doing so does not. So we might characterize the skeptic's

approach as the epistemic approach that avoids error, but it is not an approach that attains truth.

This is precisely the sense I mean when I claim that we value truth more than we care to avoid error. While both are important, if we have to choose, we will risk some error in order to attain truth, even if it is only possible truth. Not relying on sensory experience dooms us to the loss of truth *even if* we are agents in normal worlds; it puts us in the same boat as the skeptic even if we are agents in normal worlds. Such a practice cannot seriously be defended on the grounds that it is the only epistemically rational or responsible path. That agents in skeptical worlds would have less false beliefs by not relying on sensory experience is of little comfort considering that their restraint does not result in the attainment of truth. The avoidance of error becomes important specifically in contexts where avoiding a problematic method of belief formation opens the door to the attainment of truth.

In summary, if our choice is to follow the skeptic and not rely on a method of belief formation unless it has been proven to be reliable in the most general sense, or, alternately, to assume reliability unless or until there is evidence to the contrary, the latter practice lends itself better to the pursuit of truth than the former. The former lends itself to the avoidance of error, but not the pursuit of truth. This is the sense in which we cannot pursue both our primary epistemic goals fully. I contend while we are concerned about error, we care primarily for truth. This hierarchy grounds assumptions of reliability, grounds treating skeptical hypotheses as irrelevant, and grounds the claim that agents are not blameworthy for relying on sensory experience. If assumptions of

reliability are truth-conducive then there are no grounds for blaming agents for operating with such assumptions.

I have argued that our entitlement to treat skeptical alternatives as irrelevant results from our entitlement to prioritize the pursuit of truth over the avoidance of error. This ranked priority underlies assumptions of reliability and the Assumption of Reliability principle that goes as follows:

AR Principle: Agents are entitled to assume the reliability of a method of belief formation unless (1) there is evidence to the contrary or (2) there exist evidence/reasons of unreliability that a reasonable agent should be aware of.

The AR principle results in significantly more truth than its competitors. The competitors under consideration in this argument may be formulated into the following two principles: the skeptic's principle, and the reckless believer principle:

SK principle: Agents are entitled to rely on a method of belief formation *only if* the method has been proven to be fully reliable.⁵¹

Or the following looser Reckless Believer principle might serve instead:

RB principle: Agents are entitled to rely on methods of belief formation despite known or easily known evidence to the contrary.

The justification for AR lies in the truth-conduciveness of this principle over its competitors. It turns out that in light of various limitations of the first person and the facts of fallibility that the SK principle results in the conclusion that no methods of belief formation fulfill the conditions of this principle. The thought experiment in the next

⁵¹ We might formulate the skeptic's stance in a number of ways. This formulation does strike me as fair to the skeptic. In raising a relevant alternative, the skeptic is pointing out that possibility of error that seems least likely to have been accounted for. To insist that we rule it out is to demand that we have non-circular, non-question begging, evidence that shows the reliability of sensory experience, i.e. which shows that we are not brains in vats. This is what I mean when I characterize the skeptic's stance as one that requires we endorse only those methods that have been proven fully reliable.

chapter demonstrates how the SK principle fails as a reasonable meta-epistemic principle in light of these constraints. It is clear that the SK principle fails to attain truth, because it fails to endorse any method of belief formation. The Skeptic has traditionally defended the SK principle as the epistemically rational approach, but in this light we can see the skeptic's approach is best summarized as the approach that abandons the epistemic project all together (this is not surprising). The abandonment of the pursuit of truth is supposedly justified according to the skeptic because it involves some risk of error, but if it is true that a concern for error *derives from* the intrinsic value of truth, then we can see the skeptic's approach as epistemically backwards.

And, clearly, the Reckless Believer principle results in too many false beliefs. Additionally, the false beliefs that result from the RB principle are likely to cost truths as well. The RB principle is liable to cost truths for a number of reasons: reliance on a faulty method is likely to come at the cost of using a more reliable method that results in truth, and false beliefs act as filters once they become background beliefs thus increasing the likelihood of further false beliefs. So both of these alternative meta-principles result in the loss of truth; one because it is too concerned with error and the other because it is not concerned enough.

3.3: Conclusion:

The distinction between epistemic and merely pragmatic concerns has traditionally been made too sharply.⁵² The constraints in the first person are not the result

⁵² Jason Stanley gives a detailed discussion of this issue in *Knowledge and Practical Interest*. Though I do not wish to endorse a version as strong as the one he presents; I am contending the milder version that our

of prioritizing pragmatic or psychological factors over epistemic ones, as both Hume and Descartes claimed. Instead, our goal in everyday life continues to be the epistemic goal of attaining truth. It is this primary epistemic goal combined with various concrete limitations in everyday life that produces the result that if we are to attain sufficient amount of possible truths, we must ignore mere logical possibilities.

It is the concrete limitations of the context that creates the scenario in which two fundamental epistemic values are in conflict – the goal of attaining truth, and the goal of avoiding error. Thus, we must balance these two and occasionally sacrifice one for the other. Thus, in everyday life we are concerned when we have evidence that indicates our method is unreliable, or we are concerned if there is evidence of error that an agent should be aware of. However, being concerned about mere possibilities of error is to be concerned about error to the exclusion of the attainment of truth in a way that is not acceptable.

Thus, there is more to what makes an alternative relevant than the mere *mentioning* of the alternative. Such a criteria is simply too loose. If this were our criteria, belief formation could be held hostage by any possibility someone might raise. The consequences of such a loose practice would again come at the cost of not having time enough to attain possible truths.

I contend that the view put forward here accounts for a very prevalent intuition regarding skeptical hypotheses. A common response to skeptical worries comes in the

concrete parameters should be taken into account when assessing whether or not belief is justified. I am arguing strongly that parameters of belief formation must be taken into account particularly in the first person aspect of justification. Nonetheless, as will be clarified in a later chapter I do not mean this to be an endorsement of an overly subjective view of justification. I will ultimately aim to argue that our reliance on sensory experience is both subjectively and objectively justified. We possess truth-conducive justification for relying on sensory experience both from the perspective of the knower and from an objective perspective.

form of the “Aw, come on!” response DeRose noted from many of his students when presented with skeptical hypotheses.⁵³ Indeed, many find the contemplation of skeptical alternatives ridiculous even in the philosophy classroom. The prevalence of this response in everyday life in which real consequences are in the balance is even higher. This accounts for the lack of skeptical alternatives under serious consideration in everyday life. Under the view here the reason for this response is that skeptical alternatives are only logically possible and the logically possible is not relevant in everyday life because allowing such hampers the likelihood of arriving at well founded and *probable* hypotheses in the time frames given in most everyday circumstances.

When accounting for skepticism many think that epistemologists are required to provide an account of the appeal of skepticism as well as account for the possibility of knowledge. It seems more is required. Skepticism is a distinctly philosophical problem. It is surprisingly absent in the discourse of everyday life. Everyday agents think nothing of dismissing skeptical alternatives even after they have been raised. This is so to such an extreme that it is almost an embarrassment to Philosophy that we continue to invest so much time in *warding off the skeptic*. To present the conclusion to the world that we have finally found reason to dismiss the worry that *we are really brains in vats* is not likely to be met with relief on the part of non-philosophers. Most solutions to skepticism cannot account for how *very wrong* skepticism seems to most. The view here captures how very backwards the skeptic’s approach is. When we dislodge the skeptic’s “King of the Mountain”⁵⁴ stance, and evaluate skepticism as a general epistemic approach and

⁵³ DeRose, “Introduction,” 3.

⁵⁴ The “King of the Mountain” phrase is used by Keith DeRose to describe the challenge of trying to provide a refutation to skepticism. His contextualist argument relies heavily on the notion that skeptical alternatives embody possible worlds that are remote, and, thus, are not likely and not relevant. He rejects

factor in the realities of belief formation *outside* the philosophical context, we can see how crazy it seems to sacrifice the pursuit of truth simply because we *might* be wrong. When we say that it sounds odd to say “I know p, but I might be wrong about p” what we mean by that is that we have ruled out all relevant error-possibilities. Claiming knowledge with the awareness that there may be some unforeseen, unknown, unaccounted for error is not counter-intuitive. What is counter-intuitive is claiming knowledge with real and probable error possibilities that have not been ruled out. Of course, we are not endorsing knowledge ascriptions as consistent with the latter.

Many have given contextual analyses of standards for evidence. On many such analyses the standards for adequate evidence rise when the pragmatic and non-epistemic consequences are high. This sort of view accounts for the fact that our epistemic goals are not our only concern, and as non-ideal agents forming beliefs both our epistemic and non-epistemic concerns factor into the weight given each specific goal at the moment.

However, the view under consideration here is not quite of the same variety. The argument under consideration here is founded on the claim that our epistemic norms themselves can conflict in various contexts in much the same way that moral norms can conflict. In just the same way that we cannot fully value the freedom of the individual and protect individuals from outside harm, we cannot fully avoid error and pursue truth at

the notion that he is begging the question against the skeptic. However, since his position involves the claim that skeptical worlds are remote possible worlds, his argument involves an assumption about which possible world is actual (ie. He assumes we're in a normal world), and moves from there. In the argument I have given I avoid such an assumption. My argument to dismiss skeptical alternatives is grounded in the claim that skeptical alternatives have no evidence or reasoning to show that they are likely, thus they are not sufficiently grounded to warrant the time it takes to consider them seriously. Nonetheless, we might think of the argument given here as offering support for those fundamental assumptions; the assumption that sensory experience is in fact reliable in our world is just such a fundamental assumption. If such an assumption can be shown to be epistemically advantageous, then perhaps we are not begging the question against the skeptic after all. If there is a true stale-mate between the skeptic's and non-skeptic's positions, then such “small” considerations as that which is advantageous both practically and epistemically may be sufficient to justify the non-skeptic's stance.

the same time. Granted, there is much room to overlap both of these goals, but the skeptic's stance is one of being so concerned for error to the point of excluding the pursuit of truth. So it is for epistemic reasons on this view that we ignore some possibilities of error. Namely, it is for the sake of raising the probability of arriving at truth that we ignore some possibilities of error.

I think this is the heart of the "Aw, come on!" complaint against skepticism. It is not sufficient to give a solution against skepticism that gives us a technical way out from its pernicious conclusions. We need an account that not only shows us where the skeptic goes wrong, but also does justice to the intuition that the contemplation of every error possibility no matter how un-grounded is not good epistemic practice. This is the intuition behind the "Aw, Come on!" response.

As noted earlier in the chapter, there are some clear limitations to this response. The response is designed to articulate what values the skeptical and non-skeptical positions embody. As a result of the largely normative evaluation given here, the discussion is not one that will yield a refutation. Rather, it yields a discussion of the normative element that drives each position respectively in a way that mitigates the worry that the skeptic has the upper ground in her position regarding sensory experience.

Lastly it should be noted that the argument given above is, in many ways, not all that surprising. If we reject Descartes' standard of setting aside that which admits of any doubt, the standard of certainty, then we are likely to find skeptical hypotheses less than motivating grounds for doubting or rejecting our external world beliefs. The debate centers on which standard; that of fallibility or infallibility is the true standard of everyday life. I contend the extent of knowledge claims made and maintained in light of

the possibility of error, the lack of practicing skeptics walking around, and the view of science as paradigmatic knowledge attainment as evidence that the concept of knowledge as we use it in everyday life is consistent with the possibility of error. That is, we seem to endorse a view of knowledge that is consistent with fallibility and stands in stark contrast to the standard Descartes set for himself. If this is so, then the challenge rising out of skepticism becomes to understand how such fallible knowledge is possible and how we are justified in claiming such knowledge in light of our less-than-definitive grounds. In the next chapter, we will specifically evaluate our reasons for endorsing an assumption of reliability regarding sensory experience. If assumptions of reliability are epistemically justified, then so too are the beliefs grounded in them. Just as Stroud notes that if our sources of our beliefs are found wanting, then beliefs grounded in them are also found wanting.

Such a view gives a sense of how knowledge might be fallible; if the grounds we have for relying on our primary methods of belief formation are only justified *assumptions*, then the possibility of error is present at the very source of our external world beliefs.

“Justifying Knowledge Ascriptions in Everyday Life”

4.1: Characterizing the Problem

Recall the meta-principles regarding assumptions of reliability articulated in the conclusion of chapter 3; we noted three distinct possibilities: the Assumption of Reliability principle (the AR principle), the Skeptic’s principle (the SK principle), and the Reckless Believer principle (the RB principle). The AR principle is as follows:

AR Principle: Agents are justified in assuming the reliability of a method of belief formation unless (1) there is evidence to the contrary or (2) there exist evidence/reasons of unreliability that a reasonable agent should be aware of.

And the SK principle that underlies both the under-determination of the evidence and closure versions of skepticism is as follows:

SK principle: Agents are justified in relying on a method of belief formation *only if* the agent has evidence of reliability at a general level.

While we took a stab at defending the AR principle over the SK principle in the last chapter, the aim of this chapter will be to make a stronger defense of the claim that the AR principle is conducive to the attainment of truth, whereas the SK principle is not.

In order to fend off the criticism that I am equivocating on the concept of a method as truth-conducive, I must defend the sense of this concept being employed through-out this argument. Thus, I will argue in this chapter that the notion of truth-conducive as truth-generating is the best analysis of a method of belief formation in the

context of skepticism. This argument will turn on the fact that in light of skeptical worries we must evaluate reliability from a trans-global perspective, however, a notion of truth-conducive that is defined as more-truth-than-error introduces a notion of trans-global reliability that is not relevant in epistemic analyses. For this reason I will argue that we will want to analyze sensory experience from a trans-global perspective without endorsing a notion of trans-global reliability. Some of this argument will involve arguments to the effect that it is global reliability that is relevant in epistemic evaluations and that since such assessments cannot be made in light of skeptical alternatives, the trans-global perspective allows us to evaluate the AR principle in general, but it does not motivate the move towards a notion of reliability founded in trans-global reliability.

The above argument is meant to go some way toward defending the original claim that the AR principle need only be rational in order to be justified. We might formulate this claim in the following way: the skeptic insists that presuppositions or hinge propositions must be epistemically justified in order for the more local claims that depend upon them to count as justified. Alternately, our claim is that presuppositions, or hinge propositions, are not directly justification conferring, and, thus, need only be epistemically rational in order for the more specific knowledge claims to be justified. Our argument for this thus far has been that the latter standard is conducive to the pursuit of truth, whereas the former standard is so stringent as to come close if not amount to a standard of certainty regarding the evidence that underpins knowledge claims. Additionally, as argued in the last chapter, such a standard sacrifices the pursuit of truth to a concern for error to such a large extent that the standard eclipses the pursuit of truth in any context in which certainty is not guaranteed. Thus, we will see in this chapter as in

the last that it is ultimately our desire to attain truth that underpins our fundamental epistemic attitudes towards error. In the last chapter we saw that the pursuit of truth and the avoidance of error are ultimately epistemic values that drive our first person practices, and like many fundamental values the pursuit of these two epistemic goods can conflict, and, thus, standards must be developed and embraced that balance the pursuit of these two epistemic concerns. It will be argued more strongly in this chapter that AR reflects the best balance between these two epistemic concerns. Further, it will be argued and shown through a thought experiment that our concern for error is ultimately derivative of the goal for truth, and that as such it is in a sense logically backwards to sacrifice the pursuit of truth out of the possibility of error.

Thus, in this chapter, as in the last, much will be made of the logical distinction between error that is likely (i.e. error for which there exists evidence or reasons that indicate that error is actual or highly possible) versus error that is merely logically possible (i.e. error which logic tells us is possible but for which there exists no indicators that it is in fact actual error). It will be argued in this chapter using a thought experiment that assumptions regarding reliability are in fact epistemically rational, i.e. which contrary to the skeptic's claims relying upon sensory experience even without any epistemic evidence that it is in fact reliable in a global fashion is most consistent with our epistemic goals and concerns. This conclusion is contrary to the skeptical position outlined in the introduction of our discussion. In the introduction we noted that there is an implicit charge of irrationality behind many skeptical arguments, including both Humean and Cartesian skeptical arguments. In fact both Descartes and Hume make the charge of irrationality in a way that is not so implicit. Both note that what ultimately

accounts for our epistemic reliance on sensory experience is distinctly psychological rather than epistemic or rational. Thus, according to skeptics our everyday epistemic practices lack a certain rationality. So the argument goes that our higher reasoning faculties tell us that we have no reason to rely on our sensory experience, but our lower animal instincts drive us to ignore such reasoning. We see in skepticism the classic tension between the mind and the lower instincts. As discussed in the introduction, it is this specific aspect of skepticism that I am taking aim at. Instead, I would like to suggest a picture of an animal not in tension with itself but one that is behaving rationally and coherently in the context of a contingent and changing external world. The picture being given here is one that is optimistic about the consistency of fallible knowledge. It is also a picture in which man is a rational creature whose everyday epistemic practices in fact cohere well with his most fundamental epistemic values. Thus, I am suggesting that our everyday attitude towards sensory experience is one that is rational and able to stand against the most thorough scrutiny. Admittedly though, in order to see this we must evaluate the merits of the skeptical stance on its own; we cannot try to answer the skeptic's challenge because to do so is to concede the standard at which the challenge has been put. Rather than meet the skeptical challenge, it seems best to show why agents who rarely feel the need to account for skepticism prior to proceeding with their everyday knowledge claims are in fact the rational agents in the whole bunch.

Note again that the skeptic has no evidence to persuade us that we are wrong to rely on sensory experience – in fact, she has no evidence that indicates that we are brains in vats or that we are wrong in our beliefs. What she has, instead, is a logical possibility that is incompatible with our beliefs about ourselves and the world in which we live. So

the question as formulated thru out our discussion is to ask what the appropriate epistemic response is to such logical possibilities. As already seen earlier, we should not expect the epistemic rules here to be the same as they are when faced with *evidence* that our beliefs are actually *false*. A scientist that ignores a hypothesis that is a possibility, but for which there is no evidence, is not in the same boat as a scientist that ignores a possibility for which there is evidence of its truth.

Treating sensory experience as evidence that *indicates* the truth of skeptical hypotheses is awkward because of the massive deception involved in a global skeptical scenario. A skeptical hypothesis is not so much a contention that sensory experience indicates the skeptical hypothesis as it is a claim that sensory experience might not indicate an external world like the one we take ourselves to be in. So treating sensory experience as an indicator in a skeptical hypothesis is akin to treating lies as indicators of the truth. To my mind it is this awkward connection between sensory experience and skeptical hypotheses that is at the heart of what is problematic and difficult in figuring out how to respond to a skeptical hypothesis.

The concern at this point in our discussion becomes whether or not we are justified in thinking that we possess knowledge, and in ascribing knowledge to ourselves and to others in face of the admission to the skeptic that we could be brains in vats. Prior to looking at an argument that offers further defense of the assumption that the senses are reliable, let us qualify the kind of solution being offered here. In the context of everyday life, we rely on the senses with the whole-hearted belief that sensory experience is in fact reliable. More than that, both Hume and Descartes seem to be right in thinking that such reliance is deeper than one that reflects merely rational or epistemic endorsement.

Sensory experience provides the means for connecting to our world. Thus, it is the avenue of our mental, physical, emotional, and psychological connections to the people and world around us. Without sensory experience of any form we would be restricted to the confines of our own minds. Thus, it seems right to acknowledge that our connection to sensory experience is much more than epistemic endorsement. Sensory experience is more than a vehicle to truth. Its significance extends beyond the scope of the epistemic into the social and personal as well. Thus, we must acknowledge that reflection of our reliance on sensory experience from the limited perspective of our *epistemic* goals will fail to do full justice to our dependence upon sensory experience.

For these reasons any epistemic justification we come up with might seem weak in light of the depth of our reliance upon sensory experience. Additionally, it is important to note that some of what is upsetting in the possibility of being a brain in a vat extends beyond the epistemic to the deeply emotional and psychological. We cannot expect an approach that focuses on our *epistemic reliance* upon sensory experience to address all of the angst involved in the massive error articulated in skeptical hypotheses. Thus, we should anticipate that an epistemic approach to skepticism yield a response that justifies our reliance on sensory experience in an epistemic sense, but does not remove all of the worry involved in skeptical scenarios. I think this consequence of our approach is appealing because it accounts for a continued worry about skeptical scenarios even if we are able to show that our reliance on sensory experience is rational. Further, a solution to skepticism is not likely to come in the form of evidence that we are not brains-in-vats. Thus, skeptical alternatives will remain logical possibilities, and the epistemic and psychological worry involved in the mere possibility of being a brain-in-a-vat will persist

in light of any solution that might be found. An epistemic solution to skepticism has to articulate how we are justified in relying on sensory experience despite the metaphysical possibility that we are brains-in-vats.

This seeming stale-mate with the skeptic is articulated well by Moore. He characterizes the seeming stale-mate in the following way:

I agree, therefore, with that part of this argument which asserts that if I don't know now that I'm not dreaming, it follows that I don't know that I am standing up, even if both actually am and think that I am. But this first part of the argument is a consideration which cuts both ways. For, it is true, it follows that it is also true that if I do know that I am standing up, then I do know that I am not dreaming. I can therefore just as well argue: since I do know that I'm standing up, it follows that I do know that I'm not dreaming; as my opponent can argue: since you don't know that you're not dreaming, it follows that you don't know that you're standing up. The one argument is just as good as the other, unless my opponent can give better reasons for asserting that I don't know that I'm not dreaming, than I can give for asserting that I do know that I am standing up.⁵⁵

Note that while Moore believed he could gain the upper hand in the seeming stale-mate by noting his evidence in favor of his claim that he is not dreaming, we are looking to gain to the upper hand by asking what reasons there are for endorsing the AR principle over the SK principle.

Before moving on, let's remember what ground we have made in defending external world knowledge claims thus far. First, it's been established that relying on sensory experience is conducive to attaining truth in the context of everyday life in which there are time restrictions on belief formation. The result is that in such a context, the goal of attaining truth conflicts with the goal of avoiding false beliefs, with the conclusion that the two must be balanced. The result of balancing these two epistemic goals is that we are epistemically entitled to assumptions of reliability. In everyday life we work with a looser requirement than that of ruling out every single alternative.

⁵⁵ See Moore, "Certainty," in *Epistemology*, 32.

Rather, the de facto requirement seems to be that we rule out all alternatives that are sufficiently likely. This practice is epistemically defensible on the grounds that it is conducive to the attainment of truth. Note that this rule allows a view of relevant alternatives in which an agent could dismiss a skeptical alternative in a given context merely on the grounds that there are not enough evidence/reasons to think the alternative is likely. Such was the case in thought experiment G given in chapter 3. In that example it was clear that because the skeptical agent's hypothesis lacked evidence it was highly rational if not prudent for the counsel to ignore her hypothesis as time-wasting, and, therefore, destructive to the epistemic goals of the context.

Next, remember that we are not endorsing knowledge claims on grounds that justify saying "I know that p even though it's possible that not-p". Rather, the position being developed so far is that skeptical alternatives are not treated as relevant alternatives in everyday life for reasons like those given above and in previous chapters. Instead, the position, thus far, is that our knowledge claims in everyday life are the sort in which S claims to know that p appropriately in contexts in which S has ruled out all relevant alternatives.

Let us look at the reasoning behind the skeptical position that we fail to possess grounds for *claiming* to know. Such reasoning would go that in order to claim to know that p we must have a justification for *thinking* that we know that p. Such a reason will often be tantamount to our justification for our belief that p. Even if such is not the case immediately, often, upon reflection, we will provide a claim to know that p based on reflecting on our justification for our belief that p. So, the reasoning will go that in order to possess a justification for belief that p, we must have a truth-conducive reason for

believing that p. Not any old “reason” amounts to epistemic justification. For example, an agent may want to treat unreliable hearsay as reason for believing some proposition, but that does not make unreliable hearsay sufficient to ground a knowledge claim. Epistemic justification is unique in that it has a connection to truth. This connection between epistemic justification and truth is one that prevents accidentally true beliefs from counting as sufficient for knowledge.⁵⁶

However, we should be able to see upon reflection that we do not know whether or not the usual justification for our ordinary beliefs about the world is actually truth conducive and reliable or not. At least this is true in the context of skepticism when the issue is the reliability of the senses *in general*. We might have various reasons for thinking that a justification is truth conducive, and we may even possess various kinds of evidence that show at a local level that our justification for p is in fact truth conducive. However, after the skeptic raises her skeptical possibility we must recognize that we cannot *prove* that our grounds for believing that p are in fact truth-conducive. Whether or not they are will depend upon externalist facts about the world in which we live and the mechanisms by which we form our beliefs that are cognitively beyond our grasp. That is, we cannot prove that we are not brains in vats and, thereby, prove to the skeptic that sensory experience is in fact conducive to forming true beliefs. Upon recognizing this, the skeptic will argue, we need to concede that we fail know that our justification for belief p is truth conducive, and, thus, it is not appropriate to claim to know that p.

⁵⁶ I am purposely leaving the details out here. There are many ways in which we might characterize epistemic justification; whether we want to treat epistemic justification as involving some defeater clauses that prevent Gettier cases or offer some other analysis of epistemic justification is not overly important for our work here. I simply want to acknowledge that epistemic justification is of the sort that counts towards truth. Thus, in the end, if our external world knowledge claims do in fact amount to knowledge it will be largely because of the fact that in our world sensory experience is in fact reliable, and, thus, does in fact count towards the truth.

Upon reflection we can see that this reasoning is a bit too quick. The skeptic's requirements on appropriate knowledge ascription are more stringent than initially seem. The skeptic requires that we not only have a reason for believing that p, she is requiring that we have evidence that our justification *is* truth conducive in the form of having evidence *against* the skeptical hypothesis in question. She is requiring not only that we *in fact* have good grounds for our belief; she is requiring that we have *a conclusive reason* to think our justification is good as well. Of course, this will lead to the problem of the infinite regress of justification. The problem is that we must have a justification for our belief that p, and then we must have a justification for the belief that our justification is good, and so on.

What can we say in response to the above reasoning on the part of the skeptic? I think the appropriate response is to point out that such a requirement results in an infinite regress, and, therefore, is beyond the capacity of an everyday knower. Knowledge ascriptions seem to result from a belief that one has fulfilled all epistemic obligations in forming a belief that p, and, thus, has developed a stronger stance toward p than mere belief. We might offer rough alternative criteria: in order for S to be justified in claiming to know that p, she must have evidence or reasons E for believing that p that (1) are in fact truth conducive (even if this fact is external to the agent's awareness), (2) belief in p is rational and coherent in light of all of S's beliefs, (3) S has a well grounded belief that E is in fact truth-conducive. For our purposes in the context of external world skepticism (3) will amount to S's belief that sensory experience is reliable. As it turns out, S will not have any evidence that sensory experience is reliable in a perfectly general sense. It

is in this sense that skeptical alternatives threaten to undermine claims to know. It is unclear whether or not everyday agents fill criteria (3) in their usual claims to know.

The general approach I have taken so far is to treat conditions for knowledge as externalist, and to treat conditions for knowledge ascriptions as a combination of internalist and externalist requirements. An internalist type justification of the right sort seems to be a necessary requirement for knowledge ascriptions, but not for knowledge possession (at least not for all types of knowledge). However, there is nothing in this account that necessitates that knowledge ascriptions must have an internalist type justification for which there is also an internalist type justification. The requirement of internalist type justifications need not go up the entire chain of justifying reasons. As argued in a previous chapter the underlying assumption that sensory experience is reliable need not be justified with evidence of an internalist sort as the skeptic seems to require. It seems sufficient to justify this general stance towards sensory experience to note that such a stance is conducive to our epistemic goals. It then becomes the burden of the skeptic to show that reliance upon sensory experience is not conducive to truth, and accomplishing *this* seems highly unlikely. In this chapter we will continue to pursue the notion that only the specific claim under consideration requires an epistemic justification. The many presuppositions that act more generally, and higher up in the justifying chain, need only be rational assumptions. Rational assumptions are those assumptions that make sense in light of one's evidence, that are consistent with the values of the context, and for which there is no evidence/reasons to think are false. Here I would like to explicitly employ the standards we use in critical thinking and informal logic. In such arenas if one wishes to critique an opponent's presupposition it is not sufficient to merely

point out the mere fact that it is possible that your opponent's presuppositions could be false. Nowhere in the arena of the give and take of reasons and justifications is such weak criticism taken seriously. Yet, in the arena of skepticism we have allowed the skeptic to paralyze our claims to know simply because it is possible that they might be false. As admitted in the last chapter, whose standard is more appropriate is more a value-based disagreement than it is an epistemic one. Standards are largely normative, thus we cannot produce evidence that shows that our standard is best and the skeptic's inappropriate. Instead, we are aiming to show in detail what is involved in accepting the skeptic's standards. In doing so, I am aiming to justify those of us who accept looser standards and pursue external world knowledge in light of the clear fallibility of the senses.

So let us continue to characterize the issue as a question about whether or not we are epistemically justified in taking such a strong attitude towards a proposition as that of knowledge in light of skeptical hypotheses and our inability to rule them out in a definitive way. I will argue that various philosophers have given arguments that show that we are so justified. In particular, we'll look at Bonjour's argument to the effect that we are justified in preferring our ordinary beliefs about the world because they are simpler than their skeptical counterparts, and, therefore, more likely to be true. Additionally, we'll look at a highly modified version of Henderson and Horgan's argument that is essentially an epistemic Pascal's wager. This argument will be the detailed version of the one outlined at the end of chapter 3.

I will argue that both of these arguments as well as a number of other considerations provide us with reasons that make it rational for us to rely on sensory experience.

Returning to the arguments above, we might recall the conclusion that sensory experience is not an indicator for skeptical hypotheses. Notice, we must admit that sensory experience is only an indicator for our normal beliefs if skeptical hypotheses are false. So admittedly, we appear to be at a stale-mate with the skeptic. It looks as though we can only claim that sensory experience justifies our ordinary beliefs *if* skeptical hypotheses are false. For it is only if skeptical hypotheses are false that sensory experience does indicate what we take it to indicate.

The question becomes what reasons do we have for preferring our ordinary beliefs about the world in light of this seeming stale-mate? What reasons are not grounded in the belief that sensory experience indicates a normal external world? Bonjour's anti-skepticism argument starts from this stale-mate. Bonjour accepts that we cannot assume that sensory experience indicates a normal external world in light of skeptical hypotheses, and asks what grounds we have for preferring our ordinary beliefs despite this concession.

Bonjour argues that our grounds for preferring our ordinary beliefs about the world stem from the fact that our ordinary beliefs are simpler than their skeptical counterparts.⁵⁷ Skeptical hypotheses are inherently more complicated because they require a mechanism that accounts for the deception. In other words, skeptical scenarios must *explain* how the mismatch between thought and the world occurs. Such explanations may be in the form of an evil demon, a mad scientist, or sophisticated computers.

⁵⁷ See Laurence Bonjour and Ernest Sosa, *Epistemic Justification*, 95.

Whatever the explanation, the result is that skeptical scenarios are more complex than ordinary beliefs about the world. The resulting complexity stems from the fact that skeptical hypotheses must explain our sensory experience and then posit a mechanism for deception that makes it look as though we live in a normal world when in fact we do not. Whereas, normal world explanations explain sensory experience by positing that things are largely as they seem to be.

Bonjour argues that we are justified in preferring our ordinary beliefs simply on the grounds of simplicity. The added complexity of skeptical hypotheses stems from the fact that skeptical scenarios must include a mechanism as Bonjour states that “mimics the experience that we would have if the represented world were actual and we were located in it, even though neither of these things is in fact the case”.⁵⁸ It seems that Bonjour is right that skeptical hypotheses are more complicated, and I agree that such provides some justification for our preference of our ordinary beliefs. However, the contentious aspect of this argument comes in Bonjour’s attempt to show that these grounds constitute an objectively truth-conducive, an epistemic, justification for our ordinary beliefs. The connection between truth and simplicity might be too loose to provide much epistemic justification.

While I am not convinced that Bonjour’s interesting argument provides something as strong as epistemic justification, it may well justify the claim that our reliance on sensory experience is rational. If it generally is rational to opt for the simplest explanation, then it would be rational to do so in the case of skeptical hypotheses as well.

Note another difference between skeptical hypotheses and our ordinary beliefs about the world; our ordinary beliefs make an assumption of reliability whereas skeptical

⁵⁸ Laurence Bonjour et al, *Epistemic Justification*, 93.

hypotheses tend to make an assumption of deception. That skeptical hypotheses make such an assumption is precisely why skepticism can be hard to practice in everyday life. We might wonder whether or not an assumption of truth is epistemically warranted in general. Is assuming truth epistemically advantageous? Sensory experience is not the only belief forming mechanism in which an assumption of truthfulness or reliability is our de facto position. Indeed, in testimony and memory, for example, we often make prima facie assumptions of reliability. It seems that as a matter of fact we operate in most areas with an assumption of reliability unless or until there are indicators to the contrary.

The question becomes whether or not prima facie assumptions of truthfulness are epistemically warranted. Do such assumptions tend to produce more true beliefs than the alternatives? Our argument here is going to be that they do. The alternative to such assumptions is skepticism, and an overly robust skepticism may well result in less false beliefs, but is very likely to come at the cost of many true beliefs.

Thus, we might add to Bonjour's argument from simplicity the more general epistemic principle that we are generally entitled to prefer assumptions of truth over assumptions of error. We are epistemically entitled to the optimist's position. In the absence of evidence to guide us one way or the other, the goal of truth will entitle and require an assumption of reliability for all our belief forming mechanisms. When we combine the above line of thought with Bonjour's argument about the natural fit between ordinary beliefs and sensory experience and the simplicity of our ordinary belief hypotheses, we can conclude that we do indeed have epistemic reason for dismissing skeptical hypotheses.

While the above argument may seem distinct from the line of reasoning in the previous chapter, it is worth noting that the above argument might simply be another aspect of the argument given in the previous chapter. Essentially the argument thus far has been that we are epistemically justified in ruling out skeptical hypotheses because (1) they do not enjoy sufficiently high probability to be relevant in everyday life and (2) taking them seriously might save us possible error but it comes at the loss of all possible truth. If we combine (2) with the claim that while we want to avoid falsity, truth is the primary epistemic goal, we get the conclusion that the risk of false beliefs involved in relying on sensory experience is greatly outweighed by the possibility of attaining important truths about the external world.

4.2: Henderson and Horgan's Pascalian Wager

Henderson and Horgan offer a Pascal's wager type argument as justification for relying on sensory experience. We will look at a distinct version of this argument that to my mind does provide us with the conclusion that it is rational to rely on sensory experience. The thought experiment involved in this argument demonstrates how and why the SK principle is epistemically flawed. This, in turn, will provide some measure of justification for knowledge claims. It is important to note that the argument I am giving in our discussion here has been revised considerably from the original.

Henderson and Horgan give the argument in the context of a discussion on trans-global reliabilism, and it involves the contentious claim that relying on sensory experience statistically provides more truths from a trans-global perspective.⁵⁹ This claim is

⁵⁹ This argument occurred as described in a manuscript titled *Transglobal Reliabilism*. The manuscript is forthcoming with Oxford University Press and is currently titled *The Epistemological Spectrum*. The

contentious because it requires that out of all possible worlds, there are more normal worlds than there are skeptical ones. This assumption is necessary to generate their claim that relying on sensory experience is trans-globally reliable regardless of which possible world an agent is in. This claim allows them to claim that since sensory experience is trans-globally reliable, we are justified in relying on it regardless of which world we are in. However, it seems unlikely that there are more normal worlds than skeptical ones. I will be using a version of the argument that does not rely on this assumption.

Additionally, even if it were true that sensory experience is trans-globally reliable, it is not clear what motivates the conclusion that we are, therefore, justified in relying on it. We cannot ignore the relevance of whether or not the method is reliable *in our context*. Thus, I will argue that we want a different analysis all together. We need a way of analyzing the value of relying on sensory experience that does not involve a notion of reliability when spelled out as produces-more-truth-than-error. Rather, we need to see the epistemic value of relying on sensory experience despite the inherent risk in relying on sensory experience. We cannot get out of skepticism by downplaying this risk; we must articulate how and why we are entitled to rely on sensory experience despite that we might be brains in vats.

The argument given by the skeptic that we fail to possess knowledge rests on the requirement that we need to possess evidence *in favor* of our ordinary beliefs and *against* skeptical hypotheses. However, we do not possess such evidence, and, thereby

argument as it appears in the forthcoming manuscript is apparently significantly altered from the one described here, and, due to the problem of infinities, no longer contains the contentious aspect of the argument that there are more normal possible worlds than skeptical ones. However, I have not had access to this new version of the argument, and do not know how the argument has been re-characterized to yield the same result. Thus, the argument described here is the former argument occurring in the earlier version of the book.

fail to count as knowing according to the skeptic. The externalist, as an example, may avoid the skeptic's conclusion by arguing that it is not necessary that we possess any internalist type evidence in order to count as knowing, and, so, therefore need not possess and produce evidence that speaks against skeptical hypotheses. What matters for the externalist is not what evidence we can produce, but the *actual* reliability of our belief forming mechanisms. What the externalist can conclude is that if our sensory experience is in fact reliable, then we count as knowing. Yet, as we have seen such a response seems insufficient in the case of *knowledge claims*. It seems awkward to imagine an agent *claiming* knowledge on grounds that are external to the agent upon reflection.⁶⁰

Further, this sort of conditionalized conclusion about knowledge may be seen as problematic. And, indeed our conditionalized conclusion leaves the door open for the skeptic to point out that in our day to day lives we *affirm* the antecedent of our conditional. We function as though sensory experience is in fact reliable. Further, the skeptic can argue that many if not all of our claims to know rest upon the assumption that sensory experience *is* reliable. So the skeptic can argue that even if an externalist move can resist her anti-knowledge conclusion, it still leaves many if not all of our first person epistemic endeavors vulnerable to her argument.

The following response to the skeptic is best thought of as an epistemic version of Pascal's wager and it goes something like the following:

⁶⁰ In other words, even if it seems plausible that an agent know on grounds that are external to the agent, it seems implausible to say that an agent can be justified in claiming knowledge on grounds entirely external to the agent. The reason is that knowledge ascriptions are the result of conscious reflection regarding the status of one's beliefs, to justify these claims even for agents who have no reason to think they do count as knowing seems problematic. Thus, even as externalists there remains a problem for knowledge ascriptions rising out of skepticism.

1. If we are in a skeptical world, then we have no way of getting information about the contingent aspects of the external world.⁶¹
2. If we are in a non-skeptical world like the one we imagine ourselves to be in, then we will gain information about the contingent aspects of our environment via the senses.
3. Therefore, The only possible way (that we can imagine and know of) to get information about the contingent aspects of our environment is via the senses.⁶²
4. We cannot tell in a non-question begging way which world we are in.
5. Therefore, *if* we are going to understand the nature of our environment, it will be through the senses.
6. Therefore, since our aim is to achieve true beliefs about the external world, we are justified in relying upon our sensory experience *regardless* of which world we are in.

The premise claiming that we cannot determine which world we are in, premise 4, is important to the argument. Without it, the more intuitive approach to the question about whether or not relying on the senses is a good epistemic move would be to determine whether or not the senses are reliable. Premise 4 in effect states that no such determination of the reliability of the senses can be made, because no determination about which world we are actually in can be made. Premise 4 is the core of this argument. The above argument makes sense only in the context of accepting that we cannot determine which world we are in. Our discussion of epistemic perspectives allows us an *explanation* of premise 4. We are in a position to claim that such limitations

⁶¹ This follows from the fact that in any scenario in which deception is generalized, we have no way of getting accurate information about our environment.

⁶² There is the issue of semantic externalism and the potential implication that we can discern contingent facts about our environment by reflecting on the contents of our thoughts together with the thesis of semantic externalism. However, I think that this is a contentious enough issue that we need not address here. I do not think that it seriously threatens this premise. The second thing that needs to be acknowledged here is that I have explicitly restricted myself to the claim that we are talking about how to gain information about the contingent aspects of our environment. I have restricted myself in this way so as to avoid the entire issue of whether or not there is a priori knowledge and the nature of such knowledge. The final thing that needs to be mentioned in regard to this part of the argument is that rather than claim that the only possible way to get information about environment is via the senses, I have restricted myself to the claim that the only way *we might* be able to get information about the environment is through the senses. I have restricted myself in this way because we can surely imagine worlds in which agents are reliably “hooked up” to the world but they are hooked up some other way.

are inherent to our first person perspective and that the norms governing the first person need to fit the *possibilities* of the first person.

Note that shifting to the third person will not help us determine the issue in premise 4; the issue of which possible world we are actually in. Third person perspectives address issues involved in stipulating criteria and knowledge from the perspective of an outsider, and premise 4 deals directly with the question of what our specific position is. It's the difference between having a map that tells you how to get from point A to point B, and knowing where *you* currently are.

As noted, there are some differences between the argument I have given above and the argument that Horgan and Henderson present in their book. They argue that since we cannot determine which methods of belief formation are globally reliable, it is good epistemic practice to rely upon those methods of belief formation that are *trans-globally* reliable. Their argument goes that since there are more normal worlds than skeptical ones, sensory experience is trans-globally reliable. Further, since we cannot determine which methods of belief formation are globally reliable, relying on those methods that are trans-globally reliable is most likely to yield true beliefs. However, as mentioned, it is not clear mathematically that there are in fact more normal possible worlds than skeptical ones. Without this part of the argument, Henderson and Horgan have no reason to claim that sensory experience is trans-globally reliable, and, thus, no reason to think that sensory experience is likely to yield truth.

I have avoided this particular version of the argument because I think it is unnecessary to get the desired conclusion and it rests upon the claim that there are more non-skeptical possible worlds than there are skeptical possible worlds. The claim about

the trans-global reliability of the senses allows Horgan and Henderson to generate the conclusion that the senses are a trans-globally reliable method of belief formation. Thus, in any world, relying on sensory experience is epistemically justified because from the trans-global perspective the senses are known to be a truth-conducive method of belief formation. So in their version, the argument is that since there are more non-skeptical worlds than skeptical worlds, in most possible worlds, relying upon sensory experience will be a reliable way to form beliefs. Thus, we can reason that if we are aiming at true beliefs, we are justified in relying upon the senses.

An additional reason for avoiding their version of the argument is that it rests on the notion that trans-global reliability is an epistemically relevant property. However, notice that if we are concerned about attaining truth, then we must necessarily be concerned about what works in our world. It is not as though methods of belief formation that are trans-globally reliable will attain truth in our world; such will depend on the facts of our actual world. Just because a method of belief formation is not *trans-globally* reliable (perhaps it is reliable in a more local or global way) does not constitute a strike against that particular method. If a method is reliable in a given context, then it is a good epistemic tool for that context. Admittedly, if a method is very locally reliable, it might raise the worry that a small shift in context or circumstance will result in the use of a method that is no longer reliable. For example, let's imagine one lives in a town in which the clocks generally don't tell time accurately; let's imagine the clocksmith is lousy at his job. Let's also imagine that Jane who just moved to this town just so happens to own the only reliable clock in town. She uses the method of looking at her clock to tell the time. This method is reliable for her because her clock, unlike the other clocks in town, is very

reliable. Thus, her method of telling time regularly gets her to the truth. However, she lives in a context in which her method is very locally reliable. That is if she takes her method of telling time (looking at the clock) to the neighbor's house, her method will yield error. Thus, I will agree that methods that are very locally reliable can be problematic because a small shift in context raises the likelihood of error significantly. However, if Jane were aware that she had the only reliable clock, then there is nothing wrong with her method. What this shows is that very local reliability might be concerning, but there is very little chance that we will shift into another global context. Thus, there is very little reason to prefer trans-global reliability to global reliability.

Thus, Henderson and Horgan have a good case for arguing that global reliability is preferable to merely local reliability; it does not ground an analogous argument concerning global reliability versus trans-global reliability. Concerns about very local reliability motivate awareness of when and how an epistemic tool is reliable; such concerns do not motivate alternative analyses of reliability. We should not infer that trans-global reliability is better than global reliability because there is little risk of moving out of our possible world and into another one (this, of course, is quite an understatement).

Therefore, while Horgan and Henderson are right to move to a trans-global *perspective* in light of skepticism, we do not want to shift our grounds for epistemic *analysis to a notion of trans-global reliability*. What we need is to shift our perspective while retaining the *global reliability* criteria in analyzing a method of belief formation. What continues to be relevant is whether or not sensory experience is reliable in our

world, and we are concerned with whether or not we are justified in assuming reliability in the event that no such global assessment can be made in a non-question begging way.

Lastly, while it is unlikely that ordinary agents go through the reasoning involved in the Pascal's wager, it does seem that the argument reflects a leap of faith that everyday agents do make. There is a strong sense that relying on the senses and assumptions of truthfulness in general, are appropriately considered epistemic leaps of faith. While it is important to avoid false beliefs, and it is important to be mindful of error, it is also true that if we are faced with a choice between error and truth, we are well justified in opting for the attainment of truth despite a risk of error.

The skeptic's position is not one, nor can it be, that allows us to avoid error and pursue and attain truth at the same time. In fact, the latter position is exactly the position we are defending in our defense of AR. AR is mindful of error, but it is mindful of error that is sufficiently grounded. The skeptic is overly mindful of even merely possible error. The skeptic's position merely avoids error. This is what the Pascal's wager argument brings to light in a clear and explicit manner. While this is not a surprising implication of skepticism, and it may even be one the skeptic is happy to endorse, our conclusions here should reassure those who are concerned about their reliance on sensory experience in light of the skeptic's argument.

4.2.2: Grounding Assumptions of Reliability as Rational:

The conclusion of the above argument is that it is an epistemically sound practice to risk some possibility of error in order to attain possible truths. Such a practice need not be grounded in the claim that such error is unlikely. Rather, the claim is grounded in

the importance of attaining truth over avoiding error. This is particularly so in light of the fact that the error in skeptical hypotheses is unusually innocuous. This is not say that error and deception are not upsetting for a number of reasons; it is merely to point out that the error involved in skeptical scenarios *is not of the same kind* as error in regular scenarios. Error in regular scenarios is (1) the kind that comes at the cost of truth, (2) has pragmatic problems associated with it, and (3) is often the result of bad epistemic practices. None of 1-3 holds of the error found in skeptical scenarios.

Whenever possible we want to make determinations about the reliability in a local and global sense. This assures that we are using methods that are in fact reliable in our world, thereby maximizing truth. However, such determinations cannot be made with regard to skeptical scenarios. Thus, once we allow skeptical alternatives to become relevant we must make an assessment of reliability on other grounds. At such a point, we might do as above and look at the results of assuming reliability in all possible worlds.

Lastly, it is important to note that the resulting justification from the above arguments is not overly subjective. I have talked in a largely subjective manner because I have been focused on the first person perspective. Thus, I have claimed that agents may reject skeptical hypotheses because the goal of truth outweighs the avoidance of error in this particular case. However, notice that the result of our Pascal Wager argument is a justification that is actually truth conducive. *Agents that rely on sensory experience will attain more possible truths; that is they are attaining more truths out of the set of possible truths than agents who endorse the SK principle.* External world truths are not possible for agents in skeptical worlds; these truths are not relevant for them. The lesson of

skepticism is that if it is truth we seek then we must ignore some possibilities of error; we must seek truth despite the possibility of error.

The above holds because the agents in the skeptical scenarios cannot attain truth regardless of what they do. Thus, if brains in vats rely on sensory experience, it will not *decrease* the amount of truths attained in all possible worlds. If all agents in all possible worlds abstained from forming beliefs for the sake of avoiding error, then the number of true beliefs would fall dramatically (in addition to the number of false beliefs decreasing). On the other hand, if all agents (where applicable) rely on sensory experience, the number of false beliefs may increase due to the beliefs of the brains in vats, but the number of true beliefs will also increase dramatically.

Thus, we need not characterize our justification in an overly subjective way. We can conclude that if we rely on sensory experience we have a greater chance of attaining true beliefs about our world. Granted we may also have a greater chance of attaining false beliefs about our world, but the alternative to such false beliefs is not truth, the option for BIV's is to remain epistemically neutral with regard to their sensory experience – such does not enlighten the BIV to the nature of her world. If you consider that the BIV does not have any way to know if she is a BIV, we should conclude that it is epistemically irrational even for the BIV to abstain from relying on sensory experience. For all she knows, she is sacrificing truth because she might be a BIV.

If we were to create a matrix assigning value to the various possibilities in the above thought experiment it would turn out that the negative value of the false beliefs incurred should we turn out to be BIVs is less than the positive value incurred if it turns out that we are agents in a normal world. Let us spell out why it is being suggested that

these values are the rational values. The argument boils down to the claim that the pursuit of truth is worth the risk of error. Of course we could argue that all important endeavors have risk, but instead we may focus on the particulars of the thought experiment. Let's remember that the thought experiment contemplates two possible courses of action: to trust our sensory experience or not. There are two possible worlds articulated for each possible action: relying on sensory experience in a normal and in a BIV world, and not relying on sensory experience in a normal world and in a BIV world. We might boil down our thought experiment into the following possibilities:

AR principle: Relying on sensory experience (mitigated skepticism):

- Scenario 1: We are in a normal world: Knowledge gained from sensory experience
- Scenario 2: We are BIVs: many false beliefs generated by relying on sensory experience.

SK principle: Not relying on sensory experience (standard of certainty):

- Scenario 1: We are in a normal world: ALL TRUTH LOST
- Scenario 2: We are BIVs: error stemming from relying on sensory experience is avoided

I am suggesting that it is rational to assign values in such a way that the overall positive value for relying on sensory experience significantly outweighs the overall negative value for relying on sensory experience in the event that we are BIVs. Additionally, I am suggesting that the lower negative value associated with not relying on sensory experience is outweighed by the fact that the skeptical stance of not relying on sensory experience has no positive value associated with it; this stance involves the loss of truth if we are BIVs. This is the clear and explicit sense in which these two values conflict, and we have to make a choice between pursuing truth and risking error or sacrificing truth and not risking any error.

Let me say a little more why I think the loss of truth in the scenario that involves the skeptical stance is worse than the error incurred in the AR stance if it turns out that we are BIVs. The difference here lies in the distinction between the two possible worlds: the BIV world and a normal world. In a BIV world truth is lost to the agents who exist there. Such agents have not been given any method that reliably yields knowledge about their environment. The pursuit of truth in skeptical contexts is entirely frustrated by the mechanisms of deception. The negative value of error in such a context is not the same as error in a context in which truth was possible and somehow lost. The epistemic efforts of the BIV do not yield knowledge or truth no matter what the BIV does or doesn't do. She might sacrifice her epistemic agency and remain epistemically neutral with regard to all that she was not absolutely certain about (which of course will be everything), but doing so does not gain her anything. She is as lost to the nature of self and world as she ever was. The value of being epistemically cautious, the value of trying to avoid error lies in the resulting maximization of truth. The BIV world is epistemically hopeless, and it is so regardless of what the agent does or does not do. The negative value we place on the error of the BIVs beliefs results from our sense of loss about the truth. But the BIV who relies on sensory experience has not lost any truth, for truth was not possible for her. We assess the negative value of the BIVs state from our assumption that we are in a normal world and we find her situation very unfortunate. This unfortunate reality of the BIV has nothing to do with her epistemic practices. The only lesson to be learned here is that the methods we think are reliable may not be, but what of this? This does not imply that we should not rely on them. Skeptical scenarios are not like the movie the Matrix in which there was an avenue to the truth; one could take a pill and “descend down the

rabbit hole” and wind up with the knowledge of who you were and what the world was really like.

Perhaps it would help to give some examples. Let us try to imagine what upsets us about false beliefs. Imagine a scenario in which an agent has lied to her spouse for many years claiming fidelity while carrying out multiple affairs. Why would the spouse of said agent have a right to be angry, and what would he be angry about? First he may feel betrayed, let us set that specific issue aside and focus on the lies she told him. He may rightly feel manipulated by her lies in important ways. This sentiment reflects the sense that he was entitled to the truth, it was possible for him and should have been his and she intentionally and knowingly took that from him in order to manipulate his behavior and his choices. Thus, he feels that something of importance was taken from him by her for selfish reasons. Next, he may feel that the loss of truth had important pragmatic implications for his decision making; having been robbed of the truth, he stayed in a marriage that was not what he thought it was.

Notice that in skeptical scenarios the mechanism of deception insures that truth is not possible for agents in those scenarios. They can't figure out the truth regardless of what they do. But that implies that much of the negative value we associate with error is not involved in skeptical scenarios. There is no sense in which the BIV who relies on sensory experience and lands in error has been robbed of the truth, she has not lost truth by relying on sensory experience, for it was not possible for her. Second, she has no pragmatic implications associated with her error. There is no conflict between perception and belief that happens when an agent has false beliefs in a normal world. Thus, a BIVs false beliefs cohere with everything she experiences, they cohere with what in effect is

her world. This coherence is not everything but it negates a lot of what is negative about false beliefs. None of this is to say that there is no negative value associated with the error involved in relying on sensory experience in the event that we are BIVs. If truth has intrinsic positive value, then the flip side of that is the intrinsic negative value of falsity.

The purpose of taking a trans-global perspective is not to establish that sensory experience is trans-globally reliable. Rather the purpose is to evaluate assumptions of reliability versus the SK principle regarding sensory experience. From the trans-global perspective, we can see the value of AR as an epistemic principle: *it entitles those who can attain truth to attain it*. And it does no epistemic harm to those who cannot (i.e. those in skeptical worlds). Relying on sensory experience does not prevent BIV's from attaining truths.

However, it is at this point that we recognize the greater value of truth combined with the innocuous status of the brain in the vats' false beliefs. False beliefs are generally problematic because they divert us from the goal of truth, they potentially infect our other beliefs, and they have concrete, functional, problematic consequences associated with them. However, none of these apply to the BIV's false beliefs. The BIV's epistemic state is deprived because she cannot attain truth and because her beliefs are false. There is some intrinsic negative value on possessing false beliefs and it is this that defines the BIV's state as epistemically impoverished. However, much that is functionally problematic about false beliefs is not present in the BIV scenario, because the BIV's beliefs match her sensory experience. Thus, it is perfectly fine to reason that we are risking unfortunate but innocuous error to gain the possibility of attaining truth. Such is a fine Pascal's wager. And such a wager substantiates the position that relying on sensory

experience is epistemically rational rather than the irrational and irresponsible move the skeptic takes it to be.

At this point something needs to be said about the appeal of skepticism. Many find skeptical scenarios inherently disturbing and upsetting, and I am dismissing them and the proposed epistemic practices in light of them as epistemically irrational, though not obviously so. First, much that is disturbing about skeptical scenarios is articulated in the possibility that we could BIVs, or dreaming, or at the mercy of an evil demon. I have not denied such possibilities and the resulting distress seems well founded. Notice too that I have not denied the skeptic's claim that we do not know which world we are in. Rather, I have argued that skeptical alternatives are not relevant in the context of everyday life for epistemic reasons, and I have argued that even if we allow skeptical alternatives as relevant as we have done in this chapter, we can see that relying on sensory experience is still epistemically justified.

4.3: The Pyrrhonian Objection:

The Pyrrhonain skeptic might object at this point. Such objection would go along the lines that the Pyrrhonian both pursues truth and acknowledges the possibility of error. Such a position, it might be thought, avoids the choice between pursuing truth and avoiding error that I have claimed is an inherent part of everyday contexts. There may be two distinct versions of Pyrrhonian skepticism that we might consider: (1) the position in which one withholds all judgment including that of belief and (2) a position that involves withholding judgment in the form of refusing to claim that ordinary beliefs about the world amount to *knowledge*. The one position withholds judgment by

refraining from belief, the other refrains from *knowledge ascriptions*. The traditional view captured by Sextus Empiricus in *Outlines of Pyrrhonism* advocates withholding judgment as a means of attaining peace of mind. This withholding of judgment regards all attempts at pursuing truth and would seem to fit best with the first characterization of Pyrrhonist position.

The first position fits clearly the charges that I have made against the skeptic in this chapter; such a person is refusing to pursue truth in light of the possibility of error. The failure to form belief on the basis of sensory experience is the position that fits most, according to the skeptic, with the epistemic realities of our fallibility.

The second position, however, might seem like an alternative that offers the best of both the SK principle and the AR principle. If we form beliefs, then we are still engaging in the epistemic goal of pursuing and attaining truth. If we fail to claim knowledge of such beliefs, then we are acknowledging our fallibility. Such a position might be thought to be the epistemic ideal between the AR and SK principles.

However, position (2) does not cohere as well as we might initially think. Note that the AR and SK principles regard *belief formation*, rather than knowledge ascriptions. Belief is as much a subject for epistemic norms and the verdict of rationality as is knowledge. Whatever epistemic “irresponsibility” could be charged of someone who claimed her external world beliefs amount to knowledge could also be charged of someone who formed such beliefs in the first place. The AR position is involved in the kind of Pyrrhonian skepticism articulated in position (2) just as much as it is involved for the agent who claims knowledge on the grounds of endorsing AR. The question of ascribing knowledge seems to come on the heels of endorsing AR rather than the other

way around. It is not as though the agent who does not claim knowledge, but only belief fails to endorse AR on the exact same grounds as the agent who claims knowledge.

Part of the reason for the above is that the AR or the SK principles are issues regarding *belief formation*. Thus, the issue at this stage is deeper than that of knowledge ascriptions. *Beliefs* are the primary subjects of epistemic norms. If an agent fulfills the obligations of responsible belief formation and they possess reasons to believe that are accessible via reflection, there is no clear reason to deny the rationality of a knowledge claim. Whatever line of reasoning that might attack a knowledge claim will undermine belief as well. Once AR is endorsed, then belief based on the method in question is justified, and knowledge claims are simply a reflection of that endorsement. As counter-intuitive as it might seem to claim knowledge in light of the possibility of error, it is equally counter-intuitive to believe but to refuse to claim knowledge. In other words, it seems that any epistemic norms violated by AR occur at the level of belief formation, not at the level of claiming knowledge. The knowledge claims that follow on the heels of AR merely reflect the endorsement of sensory experience embodied in accepting AR in the first place. It is not as the Pyrrhonian skeptic might wish it to be- that knowledge claims are the bearers of epistemic norms, but beliefs somehow are not. Belief can be rational or irrational, responsible or irresponsible, well founded or not. Knowledge ascriptions and the practices of giving and accepting reasons in the first person, primarily center on attempts to discern well-founded versus un-founded beliefs.

There are some aspects of Pyrrhonian skepticism that we can endorse. The difference will be that the Pyrrhonian finds the following as grounds to suspend judgment, whereas I find them an inherent part of the pursuit of truth and the attainment

of knowledge. In his *“How a Pyrrhonian Skeptic Might Respond”* Peter Klein argues in defense of some of the classic Pyrrhonian arguments. Klein characterizes an argument from Sextus Empiricus on the issue of the infinite regression of justifications as “we discover... the general Pyrrhonian claim that arguments that end, end either arbitrarily or commit the fallacy of begging the question”.⁶³ Klein notes that “Either foundationalism, coherentism, or infinitism is the appropriate method of responding to the regress of reasons”.⁶⁴ The conclusion of this argument is that reason cannot settle such matters and the result is the suspension of judgment.

The argument given here can be thought of as a kind of foundationalism. The classic problem that foundationalism either lands in arbitrary assumptions or in basic and self-evident beliefs is an issue somewhat sidestepped in our approach. The latter option for a foundationalist seems appealing because it avoids the arbitrariness of the former but its’ problems stem from how such slim and few self-evident beliefs can support the whole of our knowledge. DesCartes faced just such a problem when he was faced with the deriving the whole of our knowledge from the Cogito.

If AR acts as the foundation of our knowledge based in sensory experience, then it is relatively easy to see how it can act as such a strong foundation. If we are entitled to endorse the reliability of sensory experience in a general way, then such a substantive belief might well serve as the foundation for knowledge of the external world grounded in the senses. However, the tricky question for us is whether or not AR is epistemically arbitrary.

⁶³ See Klein 2003 pg 81.

⁶⁴ See Klein 2003 pg. 78-79.

On this particular issue we are going to concede to a certain extent, but note that our position is not quite as problematic as it might first seem. First, let's note why arbitrariness is epistemically problematic. To say that a belief is arbitrary implies that one could either way on the issue; the epistemic factors do not sway one way or the other. Believing in an arbitrary manner is akin to believing on grounds that are merely lucky. If your grounds are epistemically arbitrary, then there is no more connection to the truth than if you had guessed and just so happened to get lucky. It has long been thought that knowledge excludes such kinds of luck, and that knowledge excludes arbitrariness.

If it is determined that AR is truly arbitrary then it cannot act as the kind of foundation that can confer *epistemic* justification. However, I am also going to reject the strongly internalist notion that we need to possess some kind of reasons accessible via introspection for thinking that AR is in fact true. I have explicitly argued against this criterion and characterized it as the heart of the SK principle. Nonetheless, our reasons for endorsing a meta-principle like AR should be substantive and epistemic.

Part of what has been argued so far is that our grounds for endorsing a meta-principle like AR are largely epistemic. It's just that our grounds reflect the normative element of first person practices rather than the classic requirement that justification "counts towards truth".⁶⁵ There is a sense in which endorsing AR counts towards truth—as a meta-principle it facilitates the attainment of truth. Since it is a meta-principle rather than any particular belief about the world or any state of affairs in the world, it need not be justified in such a way that our justification "counts towards truth". Rather, our justification is driven by the foundation of all epistemic endeavors; it is grounded in the

⁶⁵ This phrase is borrowed from Audi's *Structure of Justification*.

intrinsic value of truth, the pursuit of knowledge, and a reflection upon which sorts of meta-principles are most conducive to those goals.

While the skeptic notes that a general belief in the reliability of the senses underpins our particular external world beliefs, if we reflect on what underpins such a general stance towards this particular belief forming mechanism, we find a general epistemic attitude towards belief forming mechanisms in general. That general stance is reflected in the AR principle, the principle that we are epistemically entitled to assume truth, to assume reliability. It seems common sense that such an entitlement is conducive to the pursuit of truth and the attainment of knowledge. While it's true that the Achilles' heel of such an assumption is the kind of massive and general deception articulated in skeptical hypotheses, such only counts against AR if we began with the assumption that knowledge requires certainty, or that our epistemic endeavors ought to be immune from even possible error. Skeptical hypotheses have not a single shred of evidence in favor of them, they articulate the worry that we might be wrong and nothing more. While it might be upsetting to have it pointed out to us how we could be in error in a way that even the most thoughtful could never detect is psychologically upsetting, it is not the epistemic death sentence it has often been thought to be.

That our foundation is largely driven by the intrinsic value of truth is a good sign for the quasi-foundationalist view I have endorsed here. Values need not be underpinned by reasoning or justification in that classic sense of epistemic justification. Here we might recall Mill's defense of the fact that happiness is valuable; it is in fact valued. That truth is valuable is shown in the fact that we value it, and we often tend to value it for its own sake. That such a value foundation might drive our meta-principles regarding

responsible versus irresponsible epistemic behavior might not be so surprising after all. Such a foundation and such meta-principles do not stand in need of the classic sense of epistemic justification; not because we have stepped outside the realm of the epistemic, but because we have hit its' foundation.

If we were not driven by the extent to which we are by the pursuit of truth, what else would drive us, and how else would we account for the clearly normative aspect of first person epistemic endeavors? If such epistemic norms, such views about responsible belief formation did not stem from such fundamental epistemic values where would they stem from? Why should it surprise us to discover that in the real world, our fundamental epistemic values can conflict in such a way that we may not be able to fill two or more epistemic obligations in a particular context? There are notorious and well worn examples in ethics about such moral dilemmas. In fact it seems that many moral dilemmas can be characterized by a conflict between distinct values, and the conflict requires that we prioritize one value to the detriment of others. I have suggested that the skeptical problem is an example of just such a conflict in values. If this is so, then we should expect there to be no logical refutation of skepticism. Such moral dilemmas are decided and solved by choosing which value takes priority. Thus, we should expect that a solution to skepticism lies in pronouncing truth as having priority over avoiding error and defending such a choice as consistent with our most fundamental epistemic values.

4.4: A Thought Experiment: The SK principle versus the AR principle:

As previously noted, the skeptic does not think it is likely that we are BIVs anymore than the agent in thought experiment G sincerely believed aliens were planting evidence with the purpose of causing strife on Earth. The skeptic believes the force of her concerns lies in the philosophical implications for knowledge. She seems sincerely to believe that knowledge is not possible if such a claim involves any risk of error, even error that in her own mind seems far-fetched. One of the things I have suggested is that there is more to such a stance than sacrificing knowledge claims. Reliance on sensory experience can only be rational if it involves an assumption of reliability, an assumption that one is getting truth via sensory experience. To admit that sensory experience cannot ground knowledge is to admit that one ought not to rely on it. This is particularly so if knowledge claims need not involve implicit claims of absolute certainty.

It is rational to claim knowledge while implicitly stating that one does not *believe* any other alternative to hold. To claim knowledge of p and say in the same breath, “I might be wrong about p” is counter-intuitive because the claim “I might be wrong about p” is tantamount to an admission that there p is believable or plausible or likely. The agent who claims knowledge while admitting the kind of error that is merely logically possible is not involved in the latter much more contentious knowledge claim.

The question becomes, as we have framed it here, are we justified in weakly ruling out skeptical alternatives, i.e. ruling them out despite that we cannot claim to know that they do not hold. I have argued the answer to that is yes on the grounds that allowing mere logical hypotheses to hijack our knowledge is tantamount to allowing skeptical

hypotheses to hijack our pursuit of truth. Additionally, I have argued that the possibility of attaining truth outweighs the possibility of falling into error. Relying on sensory experience is more than an irrational stance driven by instinct and psychology. Rather, an assumption of reliability regarding sensory experience allows us to pursue the truth that may be available to us. The skeptic's approach is so cautious that it epistemically cripples normal agents and BIVs alike.

Let's do a quick thought experiment to drive the point home. The skeptic has used her skeptical scenarios to make vivid the possibility of relying on sensory experience only to be caught in a massive deception carried out by some evil demon or mad scientist or what have you, but what of the alternative scenarios? The skeptic herself admits that we do not know with any certainty which world we actually inhabit, but she is anxious to use her possibilities of error to encourage the most extreme caution. So before moving on, let's do a quick thought experiment of the non-skeptical variety to help us assess whether the skeptic's or the non-skeptic's overall approach to the senses is best.

Let's imagine agents who are in a perfectly normal world. By a "perfectly normal" world, I mean a world much like the ones we take ourselves to be in; a world in which the senses convey reliably about the environment in which the agents live. Let us also imagine that this world is full of agents who have been skewed by skepticism.

Now the skeptic may want to claim that such a world would be like ours with the exception that no one would claim knowledge on the basis of her senses. While I have rejected this possibility because it dwarfs this deeply epistemic issue to a minor philosophical debate about the semantics of the verb to know, I will briefly play along. On the Pyrrhonian skeptic's version our world would be one in which agents refrain from

making a judgment about whether or not what the senses tell them amounts to knowledge. First, let's be picky and point out that such restraint should come in at the level of belief itself. Thus, agents in this world ought to refrain from acting on the information provided by the senses.

However, let's grant the Pyrrhonian a break and allow the possibility that it is a coherent position to form beliefs on the basis of sensory experience while refraining from judging whether or not such beliefs amount to knowledge. Let us note that in such a world it is highly questionable whether or not the agents would survive. So first, we're going to make an evolutionary point, and then we'll make an epistemic one. What it takes to run from a bear, to fight in a war, to hunt for food, etc. is not to be found in these lukewarm convictions. I doubt the Pyrrhonian agent will have what it takes to fight a bear that she believes may or may not be there. In normal worlds, sensory experience provides a link to the outside world, and this link that keeps agents informed about outward threats. The Pyrrhonian agent that seriously entertains higher order doubts is not likely to survive a normal world. While this may seem like an overly pragmatic objection, it is unapologetically so. While epistemologists have long been inclined to distinguish the epistemic and the pragmatic, and while such distinction is theoretically useful, we are beings who cannot discern the value of the epistemic from the value of the pragmatic. That is, we are not so easily able to compartmentalize in real life as we are in theory.

Next, let us note a few epistemic things about our hypothetical Pyrrhonian skeptics. Note that our belief that the senses can ground knowledge is the foundation of every science and technology driven practice that we engage in. A world filled with

skeptics would be a world without science, such agents would not be motivated to explore, experiment, and discover in the way that we have. Such a skeptical practice would have implications in the natural sciences, health sciences, testimony, and knowledge from authority would all be impacted by the cautious reservation of the skeptic. The skeptic who does not believe knowledge is possible would not seek it out.

History is full of figures not only willing to risk the possibility of being wrong, but also willing to risk their lives for the possibility of advancing human knowledge. A world of skeptics would be devoid of the benefits of such individuals. If we imagine additionally that these skeptics do in fact live in a world where their senses are generally reliable, we must only decide that these agents are experiencing a loss that is significant and that they are doing so merely because they are not guaranteed certainty in their endeavors amounts to more than an epistemic failure, it seems to amount to a grave moral failure as well.

Surely, I will admit that when we have in mind the possibility of being a BIV, or of dreaming when we believe ourselves awake, the massive error that results is unfortunate in such a scenario. But if the question we are asking is, what is the prudent approach for the everyday knower, then we cannot look only at such skeptical scenarios. We must also consider the possibility of being normal agents in a normal world with sensory experience that is generally reliable. Further, we must decide which risk we prefer to take, do we risk the error involved in being in a skeptical world and relying on sensory experience, or do we risk being skeptical agents in a normal world? I have answered that it is, at the very least, perfectly epistemically defensible to choose to risk error rather than choose to risk truth. If this is so, then the description of the everyday

knower as one whose position is only driven by instinct is simply neither accurate nor fair. The view that the everyday knower is simply philosophically unenlightened about the error of her ways is not a fair representation of the everyday knower. Granted, it does seem that the everyday knower simply embraces sensory experience, but we ought to see that this is a defensible epistemic position upon reflection. The everyday knower is not driven by ignorance. Rather, it seems to me that she is simply driven by her pursuit of truth, and this pursuit of truth is defensible even after one has fully considered the skeptical argument.

4.5: Conclusion:

To summarize the Pascal's wager, we are either in a skeptical world or not. If we are in a skeptical world, then we have no method for attaining truths about our environment.⁶⁶ On the other hand, if we are not in a skeptical world, then we will gain knowledge of the contingent aspects of our environment through our senses. Since agents in either kind of world might ponder skepticism, they may follow the above line of reasoning to the conclusion that the attainment of truth if they are in non-skeptical worlds outweighs the possibility of false beliefs generated by relying on sensory experience in skeptical worlds.

Notice that we do want to frame the discussion in terms as black and white as those above. If what we are really worried about is how and when to trust our sensory experience, then the skeptical problem is less pernicious than the above formulation of

⁶⁶ This will hold unless the methods of belief formation or the circumstances of the scenario radically change, ie. they BIV's "wake up", or the evil demon decides to tell everyone the truth, etc.

the skeptical argument. This more specific and local skeptical worry starts from an assumption about the overall reliability of sensory experience. Notice that in order to begin with the claim that sometimes sensory experience is reliable and other times not, we must generally be trusting of the information given to us through sensation. If I found out, for example, that despite thinking there was a tree in the distance I discover, upon closer inspection, that what I thought was a tree was in fact a horse, then I have made a choice between two conflicting pieces of sensory information (x is a tree and x is a horse), and I have decided that sensory experience of seeing a horse is the more reliable information. This kind of skeptical worry is less epistemically disconcerting because it allows us to assume that sensory experience is generally reliable and requires that we be cautious and develop a sense of the parameters of the reliability of sensory experience. Thus, we have in general developed a sense of such limitations on the efficacy of perception. If a prosecuting attorney puts an eyewitness on the stand that was drunk at the time of allegedly witnessing the crime in question, she is likely not to sway the jury as much as she might like to. Such an eyewitness is not particularly credible because we generally have found sensory experience to be less than reliable when one is under the influence of mind altering substances.

Again, we have an argument founded the view that the attainment of truth outweighs the possibility of error. This is particularly so for an alternative and error that is merely *possible* error. The argument in chapter 3 was that the time constraints in everyday contexts are such that the attainment of truth and the avoidance of error conflict in such a way that we must choose which goals have priority; we must choose which norm to deviate from. The stalemate that Moore characterized has been reconceived as a

problem of conflicting epistemic norms. The skeptic represents the position of preserving the goal of avoiding error and in doing so sacrifices her pursuit of truth; even the Pyrrhonian skeptic who refrains from judgment has given up the epistemic game of trying to attain truth. The reckless believer position characterized in chapter 3 represents the position of pursuing truth while not paying enough mind to concerns about avoiding error. The assumption of reliability position is the middle ground it grants us the right to pursue truth despite that it may not be available to us if we in fact are BIVs. However, the AR position also recognizes the negative value of error by requiring that we pay attention to error that is likely, i.e. error for which there exists evidence or reasons to believe is likely or plausible.

The scope of this argument is such that it will justify assumptions of reliability for any method of belief formation if the following conditions are met: (1) there are no local indicators that the method is unreliable, and (2) the error possibilities articulated are merely logically possible. That the above argument will justify such methods follows from the fact that for such methods, it will result that the methods are reliable in some possible worlds and not others. The reasoning that we are justified in attempting to attain truth will kick in at this point. Thus, for all methods of belief formation if there are no indicators that the method is *unreliable*, then we are justified in relying on it. This argument will pertain to all methods of belief formation regarding skeptical alternatives in which error is occurring without any means for detecting, verifying, or refuting the possible error.

While I am not particularly fond of evolutionary arguments because they tend to

apply one and all – let’s go ahead and briefly note the advantage of AR from an evolutionary perspective. From a practical and survival oriented point of view it would not serve well to endorse the SK principle. Refraining from relying on sensory experience because it *might be* faulty would result in very bad consequences for those agents who are in normal worlds. Admittedly, it does not even seem rational for BIV’s to refrain from relying on sensory experience given that they too have no idea which world is actual. But this merely shows how strong the justification is for relying on sensory experience.

Reliance on sensory experience is far short of the epistemic foul the skeptic claims it to be. It is only a foul if we endorse the idea that any possibility of error is inconsistent with the possession of knowledge. We should remember that entertaining skeptical scenarios in everyday life is as time-wasting and destructive to our epistemic goals as entertaining the hypothesis that aliens are trying to create strife on Earth in thought experiment G. Many share this intuition about the relevance of skepticism, but often they are not the sort to enter the ranks in philosophy. It is worth drawing out and remembering the grounds for such attitudes towards skepticism.

Chapter 3 aimed to show that the issue of skepticism boils down to a choice between avoiding all error, and pursuing truth. In this light, it seems defensible to pursue truth. Admittedly, the stalemate remains. The skeptic can charge that in endorsing AR we open ourselves to the possibility of error. We, however, can charge in turn that endorsing SK the skeptic gives up the pursuit of truth. Rather than solve this stalemate, I have sought to disarm the skeptic’s criticism that reliance upon sensory experience without *evidence* of its reliability is epistemically irrational and unjustified.

This chapter aimed to show something more substantive; it aimed to show why the AR principle is epistemically preferable to the SK principle. In this chapter, the version of Henderson and Horgan's argument we discussed shows how endorsing the AR principle allows those for whom truth is attainable to pursue and attain those truths, whereas the SK principle epistemically handicaps one and all regardless of circumstance. And it does so in the name of avoiding error. I have argued that as a result the SK principle cannot be defended as the only epistemically responsible approach as might be argued by the skeptic.

“Skepticism and Shifting Epistemic Perspectives”

Chapter 5.1: Introduction

What emerged in the last chapter is that we may block skepticism by noting that there are relevant distinctions between the scenarios in which the Rule of Exclusion legitimately applies and skeptical scenarios. One of the biggest differences we noted between usual cases involving the Rule of Exclusion and skeptical alternatives is that skeptical alternatives are merely logically possible. Even the skeptic herself has no reason or evidence to think her hypothesis true. This difference is significant enough that it warrants treating skeptical alternatives differently than one would be required to treat alternatives raised on significantly stronger epistemic grounds.

If we acknowledge that fallibilist knowledge is possible, then we admit that some possibility of error is consistent with knowing. Embracing fallibilism will involve balancing the pursuit of truth with the possibility of error without sacrificing either epistemic good entirely to the other. Is it possible to convince the skeptic herself that knowledge is consistent with some possibility of error? Probably not. Our aim here though is to defend the epistemic practices of the everyday knower. And the everyday knower is entitled to reject skeptical hypotheses for the exact same reasons the agents in thought experiment G reject skeptical hypotheses; they are not sufficiently likely to justify time spent on them. We have all heard the phrase “time is money”, well it turns

out that time spent on error scenarios comes at the cost of something else, and, therefore, the common sense driven and prudent knower will be wary of dallying time on the merely logically possible. Might she dream about and contemplate the merely logically possible in the Philosophy classroom? Certainly. However, that does not mean doing so is prudent in real life. In the Philosophy Classroom we allow that which is not practical in everyday life; we set aside time constraints with the specific purpose of exploring issues of all stripes with the kind of depth our normal lives do not afford us. If the skeptic wants skeptical hypotheses to be taken seriously beyond the realm of Science Fiction or Philosophy, then she must show that we have real and legitimate reason to be concerned and worried about the error scenarios she frets over. Thus, the burden of proof rightly lies in her hands and the standards for raising error scenarios in everyday life are not as loose as she imagines them to be.

In addition to larger questions regarding epistemic theory, a number of concrete questions raised in the last chapter need answering in this one. Specifically, the following issues need addressing: the issue of warranted assertability, the issue of closure in relation to the view being developed here, and the details of a relevant alternatives position as articulated in chapter 3.

5.2: The Bigger Picture: How might a Theory of Knowledge Look?

Skepticism has shown us that we need to accord justifications occurring in the first person perspective a more limited epistemic role than they have historically been given. Traditionally, our reasons for our beliefs and the evidence that we have that leads us to claim knowledge has been expected to be more tightly connected with truth than it

seems is actually the case. On the Cartesian view of justification, one's evidence for a given belief needed to imply the truth of that belief in order for it to count as knowledge. And, yet, when we look at skeptical scenarios, such as the dreaming hypothesis, our evidence does not support our ordinary beliefs over their skeptical counterparts to any degree. Unfortunately, the problem is not as simple as the worry that our evidence does not rule out skeptical scenarios with certainty, because our evidence does not rule them out at all. Thus, skeptics such as Hume have argued that our reasons for preferring our ordinary beliefs over the skeptical hypotheses must be non-epistemic. If they were conducive to truth, then we would have evidence that spoke against skeptical hypotheses to at least some degree.

Still, it seems intuitive to say that people can count as knowing even if they are not certain about what it is they claim to know. If, for example, someone forms a belief on the basis of sensory experience that is in fact reliable, it is not implausible to claim that such a belief amounts to knowledge. Admittedly, both skeptics and internalists might disagree with such a claim if the agent lacks any reason to think that her sensory experience is reliable. We know, additionally, that the skeptic will claim that none of us has a general and justifying reason for relying on sensory experience. We noted in an earlier chapter that this intuition to accord knowledge despite fallibility is mistaken according to the skeptic because it requires that we have some degree of evidence in favor of our belief (over the skeptical alternatives).

What we need is a sense of justification that allows evidence to have epistemic value and weight despite that it does not rule out every alternative. I will urge that the limitations of evidence in the first person perspective stem from the inherent limitations

of the first person perspective itself. The relevant alternatives position is an account of knowledge upon which fallibility is cashed out in terms of knowing that *p* as opposed to a set of relevant alternatives, whereas an infallibilist account of knowledge might characterize knowledge as ruling out each and every error possibility.

Since the skeptic employs what we have been calling the Rule of Exclusion, the skeptic gets her conclusion by attacking our evidence. What the skeptic gets via the Rule of exclusion is that we need some *reason* to prefer our ordinary beliefs to her skeptical alternatives (this is generally the case in our everyday epistemic practices). So a tempting response is to claim that the skeptic requires certainty, but there is no clear assessment of where she deviates from our everyday epistemic practices, and it will follow that there is no clear sense about which of her premises is flawed.

If we abandon the idea that internalist justifications are the *primary* mark of knowledge, then we may be left with a way to accord them epistemic value while at the same time conceding to the skeptic that they do not, indeed, speak against skeptical scenarios. Such a view will account for the fallibility inherent in the first person perspective by noting the de facto *limitations* of the first person. Indeed, it would make sense if the justifications that we operate within the first person are limited in just the same way in which our first person perspective *itself* is limited. Our project here is to understand in terms of epistemological theory how knowledge is possible in light of such limitations.

Let me begin by noting that the view I have in mind is largely externalist in nature. The view is externalist in the sense that internalist type justifications are not required in order to count as knowing. My reasons for this are independent of the

discussion here and it seems unnecessary and not relevant to our specific discussion to defend externalism generally here. It may not be essential to my solution that a theory of knowledge developed around the idea of epistemic perspectives be externalist. It seems plausible that internalist views that incorporated a notion of epistemic perspectives could be developed as well. I suspect that the view I will develop here fits my solution in the most ideal way, but it seems somewhat irrelevant to our discussion here to try and defend this view. Instead, my aim here is to fill in the notion of epistemic perspectives in a way that allows us to explore some of the claims made in the previous chapter. Part of the issue is that any theory that acknowledges epistemic perspectives in Bonjour's sense will incorporate both an internalist and externalist aspect. The debate, then, will be centered on whether or not knowledge requires an internalist justification. So the view being developed here is one that denies the necessity of internalist justifications for the possession of knowledge. However, internalist justifications are involved in the first person perspective and, thus, will be *necessary for knowledge ascriptions*.

Since a first person approach involves concerns about whether or not our beliefs are justified, theories that focus on the first person perspective tend to offer guidelines and norms about what it is for beliefs to be justified in the first person. Both perspectives are unavoidable and important; the result is an epistemology that credits both perspectives as relevant, and the kinds of justification stemming from them important. Additionally, it will seem a brute fact that many externalist criteria for knowledge are more important than the first person conditions for knowledge ascriptions. For example, one criterion essential for knowledge based on sensory experience we can see is that sensory experience is a reliable way of forming beliefs in our world. Or we might phrase

the condition as the requirement that if we are going to attain knowledge via sensory experience, we must be “hooked up” to our world in the appropriate way.

What we need now is some sense of how internalist type justifications function epistemically in light of the various externalist elements involved in knowledge. To put it most intuitively, the view we need is one in which internalist type justifications function against the background of the reliability of other various means of belief formation. In this instance, we will be looking at how sensory experience functions as an internalist justification against the background of the reliability of the senses. It may be noticed already that this view has great similarities to others put forward. In Pritchard’s *Epistemic Luck* he talks about Wittgenstein’s use of the notion of hinge propositions; propositions that function as necessary and given within a particular context.⁶⁷ He sees various propositions about the reliability of the senses functioning as hinge propositions in most everyday contexts.

In everyday contexts, asserting knowledge is often just fine according to Pritchard. However, in skeptical contexts concerns about skeptical alternatives have made it so that various assumptions about the reliability of the senses are no longer allowed, and thus it may be inappropriate to *assert* knowledge.

The idea I’ll be outlining is somewhat akin to the above view except that it is meant to also offer an *explanation* as to why certain propositions about the reliability of the senses are hinge propositions. Sensory experience and the reliability of the senses is just an example of this. In order for sensory experience to have any epistemic value it must be *against* or *in the context of* the reliability of the senses. So propositions or assumptions about the reliability of the senses have an epistemic priority. Because they

⁶⁷ See Pritchard’s *Epistemic Luck*, p. 226-229.

have epistemic priority we cannot call them into question when discussing issues involving sensory experience.

Once we call the reliability of the senses into question, we are at a loss for a way to resolve the debate. This is why, unlike most everyday cases in which people put forward alternatives to be considered, the skeptic has no evidence for her hypothesis. She *cannot* have evidence for her alternative; indeed, she has no reason to even *suspect* that it might be true. In such contexts, we have hit the point that the first person perspective cannot go beyond.

What the skeptic has pointed out is not that we fail to know what we think we know. Indeed, she is mistaken to infer that from her argument. What she *has* gotten a hold on, though, is the limits of the first person and the gap between first person justifications and truth. Those essentially externalist requirements cannot be addressed in the first person perspective because everything that occurs in the first person is dependent upon those inherently externalist criteria. On the traditional internalist approach to knowledge, internalist type justifications were taken to be a stronger marker for truth than they actually are.

. If we really are agents who are connected up to our environment via our senses such that we fairly reliably get information about our environment through our senses, then we possess knowledge of the external world, and will often possess internalist type justifications that will give us good reason to think we possess knowledge of the world. That our sensory experience, characterized in a subjective and first person way, does not speak against skeptical hypotheses does not detract from the fact that agents in such a scenario have internally justified beliefs and are warranted in claiming such knowledge.

So our response is to own much of what the skeptic argues, but to deny her conclusions on the basis of the argument that she is mistaken about the role of sensory experience. We can see that internalist justifications have a vital but limited epistemic value because they are not connected that closely with the truth. But it is *not* arbitrary which propositions function as essentially externalist criteria for knowledge, nor should we concede to the skeptic that we should not *assert* that we possess knowledge of the external world. Though the standards for the possession of knowledge might be high enough to require that we in fact formed our beliefs in a reliable way and that they are in fact true, it is possible for us to allow that knowledge *ascriptions* reflect the first person perspective in which we are largely attempting to determine whether or not the criteria for knowledge possession has been met. First person justifications are our *measure* of whether the standard has been met, not the standard itself.

The last point is worth drawing out a bit. I urge that an externalist solution to skepticism that is along the lines of Duncan Pritchard's in his *Epistemic Luck* should not concede to the skeptic that it is inappropriate to assert knowledge of the external world. Rather, our knowledge claims are grounded in (1) our justified assumption of the reliability of the senses and (2) the relevant sensory experience.⁶⁸

Indeed, it is my sense that this is why, when confronted with skepticism, many respond with the feeling that the skeptic has performed some kind of trick and that she is almost ridiculous to demand that our evidence rule out her alternative hypotheses. Our appropriate response about knowledge assertion is that if our first person justifications indicate knowledge to the best of our ability it is appropriate to assert knowledge. This

⁶⁸ See Pritchard, *Epistemic Luck*, 227.

is so even if the skeptic points out that our evidence, and, thereby, our assertions of knowledge do not *imply* knowledge.

Let us remember that the skeptic's principle regarding reliability amounts to an unwillingness to pursue knowledge in light of the possibility of error. Contemporary paradigms of knowledge are fraught with fallibility and possibilities of error with much concrete success, and, yet, we philosophers have struggled to accommodate such views in theories of knowledge. We continue to feel overly persuaded by the skeptic that without some proof of reliability regarding sensory experience we fail to be justified in relying on it. The fundamental working assumption in this argument is that such is simply not the case and that it is worth bearing out such views.

If we follow Pritchard's argument in *Epistemic Luck*, we ought to conclude that it is not right to assert knowledge in skeptical contexts because, doing such will often generate the conversational implicature that we have evidence that speaks against rival skeptical scenarios.⁶⁹ But this is precisely the move that is mistaken. While this rule often does hold in our everyday practices, it does not hold in skeptical scenarios. Even if we were agents in a normal world, we would not possess *that* kind of evidence. Yet, *this does not imply either that we fail to know (given an externalist account of knowledge) or that our sensory experience has no epistemic value.* To see that the latter is so, again, reflect on the fact that even the sensory experience of agents in a "normal" world will not speak against skeptical alternatives and, yet, their sensory experience does indeed reliably indicate the truth, and does so *despite* that it fails to speak against skeptical alternatives. Knowledge ascriptions should only incorporate the standard that we count as knowing in

⁶⁹ This is what Duncan Pritchard claims in his discussion "Epistemic Angst" from his *Epistemic Luck*.

light of our evidence and our background beliefs. Thus, we might formulate the truth conditions for knowledge ascriptions as follows:

S's claim to know that p is true iff:

1. S does in fact know that p.
2. S has reasons or evidence to believe that she has met the conditions to count as knowing that p.

Alternately, we might also be interested in characterizing warranted assertion as reflecting the conditions in which someone is *justified* in claiming to know that p:

S is *justified* in claiming to know that p iff:

1. S has reasons or evidence to believe that she met the conditions to count as knowing that p.

Skepticism can be seen as a criticism that agents fail to meet the conditions for justified or warranted assertion. The skeptic wants to argue that in light of skeptical alternatives we cannot continue to see ourselves as meeting the requirements to count as knowing that p. Somewhere in the skeptic's beliefs is the assumption that in order to count as knowing one must be able to definitively rule out all error possibilities. Thus, there is one more way to think about the conditions of warranted assertability in the relevant alternatives camp:

S is warranted in claiming knowledge of p provided that she has ruled out all relevant alternatives to p.

The limitations of first person justifications do not prevent us from knowing.

Skepticism does not even imply that sensory experience has no epistemic value. Agents in normal worlds, who are hooked up to their environment in the right way, are gaining a good deal of information from their sensory experience. Indeed, they are able to tell that they have hands; they can tell when they are in the presence of trees, chairs, where to

step, when they might fall, etc. They achieve a good deal of truth via their sensory experience. Thus, it has epistemic value provided that one is in a normal world.

The mistake that the skeptic makes is that she demands that we be able to tell from our first person perspective that we are hooked up properly. If we cannot tell from the first person, then we cannot rule out her hypotheses. If we cannot tell from the first person which possible world we are in (a skeptical or non-skeptical world) then our sensory experience, as characterized subjectively and used as an internalist type justification, is not sufficiently connected to truth according to the skeptic. When we fail to acknowledge the various external requirements for the possibility of knowledge then it falls to sensory experience (when used as internalist type justification) to show things that it cannot show.

Given these considerations, our diagnosis of the various intuitions regarding skepticism would go as follows. I have long noticed that people often have conflicting intuitions regarding skepticism. There is a strong sense when presented with a skeptical argument that there is something about it that is undeniably correct. Namely, the fact that our evidence is consistent with whatever skeptical alternative is mentioned. When this is combined with the general concession that we usually need to be able to rule out competing alternatives in order to be justified in accepting our preferred hypothesis there is a strong sympathy toward the skeptic's argument. However, I have also noticed that despite these two concessions there is often a very strong reluctance to accept the skeptic's conclusion. This reluctance is often accompanied by the feeling that the skeptic has performed some kind of parlor trick and that she has made a mistake somewhere but that we cannot tell where.

What we can say at this point is that the skeptic's argument is compelling because she employs rules that usually apply in the context of warranted assertion. In order to be justified in asserting knowledge, one generally has to rule out all competitors. Because these rules are correct and do apply in the first person, we feel compelled to follow them and are sympathetic to their employment.

The skeptic's own lack of evidence or reason for thinking that her possibility is a genuine possibility is independent evidence that something is not right in this dialectic. While it is the case that a person is generally required to have an epistemic reason for preferring one alternative over another in order to be epistemically justified in doing so, it is also the case that it is generally true that the one proposing the alternatives to be addressed has *at least some reason* for counting the contending alternative as one that should be taken seriously. Hence, this entire dialectic starts with the skeptic's foul and ends with what seems to be a foul of our own. However, I have urged that our lack of evidence is merely a symptom of the skeptic's own foul. Indeed, we might think of the skeptic as committing a fallacy of accident. The accidental features of skeptical alternatives are such that the Rule of Exclusion does not apply in this instance. To endorse the Rule of Exclusion for skeptical hypotheses (1) conflicts with our usual response to alternatives deemed insufficiently grounded, and (2) amounts to endorsing a requirement of absolute certainty for knowledge.

We might point out to the skeptic that our reasons and our evidence are best treated as indicators and that we have strong epistemic reasons for prima facie reliance on sensory experience. Such strong prima facie warrant is not dislodged by the mere possibility of error. In this dialectic we can see that what the skeptic really wants is

certainty. She does not demand that knowledge imply truth. Her requirements are significantly stronger than that; she requires that our grounds for our beliefs show definitively that we do indeed possess knowledge. She requires that our grounds for knowledge assertion imply truth and rule out all possibilities of error. She requires not only that knowledge imply truth, but that knowledge ascriptions do as well. In this light we can see how stringent the skeptic's requirements are, and we have pushed the skeptic's argument into that earlier formulation in which the skeptic simply demands and stipulates that knowledge, and, thereby, knowledge ascriptions require certainty. There are no clear grounds for accepting this standard. On this standard what it means to rule out an alternative is that an agent knows definitively that the alternative does not hold, and, yet we do not have clear reasons for endorsing this view. In the next section, we will look at alternative criteria for what it means to rule out an alternative.

5.3: Relevant Alternatives Conditions and “Ruling Out”:

The argument in chapter 3 presents different criteria for what it takes to “rule out” an alternative. What is often considered ruling out an alternative is to count as having good evidence against that alternative.⁷⁰ It should be made explicit that I have endorsed a different concept for what it means to rule out an alternative. On the view endorsed by Smith, for example, to rule out a skeptical alternative we must possess “good” evidence against the incompatible alternative. I have not argued that we have such, nor does such an argument seem likely. Thus, relevant alternativists must either endorse different criteria for ruling out an alternative or be largely concessive to the skeptic. In chapter 3 I

⁷⁰ Joshua Allen Smith endorses such a view in his “Relevant Possibilities,” *Philos Stud* (2008) 138: 55-71.

argued that the skeptic's alternatives should not count as relevant because treating them as such detracts significantly from the pursuit of truth in a way that seems unacceptable. Thus, chapter 3 largely aimed to establish the conclusion that we are justified in treating skeptical alternatives as irrelevant in the context of everyday life. However, such an argument only gets us so far given that there are clear contexts in which skeptical alternatives are treated as relevant. Since I do not wish to argue that such contexts violate epistemic norms and that the consideration of skeptical alternatives is never appropriate, we will want a sense of how to handle skeptical alternatives in contexts in which they are treated as relevant.

Given the arguments in chapter 3 the contexts in which skepticism can be taken seriously are not merely those contexts in which someone has *merely mentioned* a skeptical alternative. Rather, it also needs to be the case that the context is such that the contemplation of a skeptical alternative is not destructive to the immediate epistemic goals of the context. The agent in thought experiment G that proposes a mere logical possibility to explain the activities of a rival nation has acted in such a way that the raising of that hypothesis is destructive to the epistemic goal at hand; namely, the goal of arriving at a well-founded hypothesis in the time allowed. So contexts in which skeptical alternatives can be relevant are likely to be philosophical contexts. Philosophical contexts are distinct in that we purposely set aside time to reflect further on our beliefs and our grounds for our beliefs. Thus, it may not be surprising that in philosophical contexts the contemplation of skeptical hypotheses does not disrupt the immediate epistemic goals of the context because the goal is to engage in deep inquiry.

Once a skeptical hypothesis becomes relevant; if we apply a notion of ruling out that requires evidence against the alternative we may well be at a loss for how to rule out a skeptical alternative. There are two things to say about this situation: (1) it seems that when we move to a philosophical context we explicitly employ higher standards for adequate evidence than are employed or are practical in everyday life, and (2) it remains even in a philosophical context that “ruling out” need not involve evidence against the hypothesis. The criteria that in order to rule out an alternative one must have good evidence against it is tantamount to the criteria that in order to rule out an alternative one must be justified in claiming to know that the alternative does not hold. If one has good evidence against the alternative, then one would be justified in claiming to know not-H (where H is some incompatible alternative). The worry here is that such a standard implicitly endorses or comes close to endorsing a standard of certainty for knowledge. This is of concern not only because such a view will wind up in skepticism, but it is also worrying because such a view fails to account for knowledge as consistent with the possibility of error. As noted earlier, it is a challenge of contemporary epistemology to develop accounts of knowledge that validate and legitimate current paradigms for knowledge.

Next, it is worth noting that this particular standard for ruling out an alternative does not seem consistent with actual practices regarding competing alternatives. Unless we wish to undermine all such practices, we had best re-think this standard. It seems that our actual practice is to treat the criteria for ruling out a hypothesis as on par with the grounds for raising the hypothesis. The better the grounds are for an opponent’s theory or hypothesis or criticism, the higher the standard for meeting the criticism or rejecting

the alternate theory. Likewise, when someone raises an alternative to our current belief; the better the grounds for raising the alternative, the more that is required in ruling it out.

It is not clear what the intuition or motivation is for thinking that a hypothesis that is not well grounded needs to be ruled out with “good evidence”. The intuition behind such a view might be that knowledge is such that it should speak against all error possibilities, even those that are not well founded. However, in this light it seems that our worry that such a standard is coming close to insisting on a standard of certainty for knowledge is well founded. If possessing knowledge means that one can rule out all error possibilities, then we are coming close to a view of knowledge that requires certainty. However, as has been noted such a view of knowledge seems inconsistent with our actual practices and inconsistent with such paradigmatic knowledge as scientific knowledge.

Instead, I am proposing a criterion for ruling out that has more intuitive appeal. It seems to me that knowledge claims are consistent with acknowledging but not endorsing error possibilities. I have granted that it does seem inconsistent to say “I know that p but it could be that q” (where q is an incompatible hypothesis). This is not what I am endorsing. On my view agents assert knowledge with the belief that all error possibilities have been sufficiently ruled out. However, the grounds for ruling out some of the error possibilities are not definitive. Thus, on such a view an agent would claim to know, but acknowledges that it is possible that some error possibility is in fact actual. The difference between the two is that on the first reading of the fallibility of knowledge the agent has no grounds for thinking that q does not hold. On my view I allow agents to rule out alternatives even if the grounds for ruling out the alternative are not sufficient to

ground a knowledge claim. Many knowledge claims that scientist make occur against the background of alternatives that have not been “strongly” ruled out. Thus, they could not appropriately claim to know that the alternative does not hold.

Given these considerations we may be inclined to endorse a view of what it means to “rule out” an alternative in which what it takes to rule out an alternative is on par with the grounds for raising the alternative in the first place. We might formulate this idea into the following criteria:

S counts as ruling out an alternative H iff:

S’s reasons or evidence against H are epistemically stronger than the grounds or evidence in favor of H.

I think that the above concept of ruling out is appealing not only for its’ role in an anti-skepticism epistemology, it is also intuitively appealing on its own grounds. The practice of giving and accepting justifications is one that is a give and take process. We should not conceive of a competing alternative as existing in some de-contextualized manner. Allowing a view in which there are no conditions for raising an alternative beyond the mere mentioning of the alternative, but endorsing a rigid standard for ruling out alternatives stacks the deck against those making knowledge claims in a way that does not seem fair or intuitive. Such a view does not seem consistent with our actual practices, and clearly stacks the deck in favor of the skeptic. An alternative for which there is not a shred of evidence, such as skeptical alternatives, can be ruled out on grounds that do not warrant a knowledge claim. This intuition accounts for why many quickly dismiss skeptical alternatives. Skeptical alternatives are mere logical possibilities; they reflect possible error, not likely error. If we think of knowledge as consistent with fallibility then this should be reflected in allowing ruling out to occur on grounds that are indeed

epistemic but not sufficient to warrant a knowledge claim that not-H. This reflects the fallibility of knowledge; the existence of error possibilities that we do not believe to hold, but might in fact be actual.

5.4: Lewis's Rule of Actuality and BIV's:

Attempting to address criticism that relevant alternative theories are vague, Lewis offers a number of criteria for when an alternative may not properly be ignored.⁷¹ One such criterion is the criterion that the actual may never properly be ignored. This criterion is highly intuitive and it should be noted that I am not endorsing a view in which this criterion is violated. However, it should also be noted that in the context of skepticism this criterion has to be formulated carefully otherwise it will turn out that skeptical scenarios are never properly ignored. This will particularly hold if we combine his rule of actuality with his rule of resemblance which holds in the event of one possibility saliently resembling another, in such a case Lewis maintains “if one of them may not be properly ignored, neither may the other.”⁷² Lewis himself acknowledges a problem with the rule of actuality and resemblance; when taken together they effectively imply that no skeptical hypothesis is properly ignored. Lewis admits this is problematic when he states, where W is some alternative:

Plainly, we dare not apply the Rule of Actuality and Resemblance to conclude that any such W is a relevant alternative – that would be capitulation to skepticism. The Rule of Resemblance was never meant to apply to this resemblance! We seem to have an ad hoc exception to the

⁷¹ See Lewis, “Elusive Knowledge,” in *Skepticism*, 227-232

⁷² See Lewis “Elusive Knowledge” in *Skepticism*, 227-232.

Rule, though one that makes good sense in view of the function of attributions of knowledge.⁷³

The argument we have put forward before may be able to help make some sense of the issues here. First though we must take note that Lewis is articulating a relevant alternatives position that acts as conditions on knowledge, whereas we are concerned with relevant alternatives as a straightforward way to characterize knowledge ascriptions in the context of skepticism. Nonetheless, we might be able to avoid some of the ad-hocery in Lewis's account if we maintain that all alternatives that are merely logical possible may properly be ignored.

The justification for ignoring the logically possible lies in embracing a view of fallible knowledge. On such a view, those claiming knowledge need not address *every single error possibility* because such a requirement is tantamount to requiring certainty in order to be justified in claiming knowledge. Instead, those claiming knowledge need to have grounds for rejecting plausible or likely error scenarios. On this view, those claiming knowledge must have reasons to believe they have avoided what we take to be likely error; error that there is some reason to believe may hold.

We might say something like the following in light of such considerations: what we want is not resemblance of any sort, but a notion of the right kind of resemblance. Scenarios that resemble actuality, or what we take to be actual, need to be well founded in order to be relevant. If not, then we risk setting a standard of certainty for knowledge. This is a standard we are blatantly trying to avoid in exchange for justifying a fallibilist approach as a rational and coherent approach to knowledge.

⁷³ See Lewis page 228.

On such a fallible-knowledge view, skeptical scenarios for which there is some reason to believe hold will count as relevant, whereas those for which there is no evidence or reasons will not count as relevant. Thus, the well-worn Fake Barn scenario would turn out to be relevant provided that there is some evidence to indicate that fake barns are likely in the area. However, let us imagine that fake barns are common in the area but that the agents in the example have absolutely no reason to think so. In such a scenario the possibility that the “barns” they are looking at are only barn facades is not a scenario that the agents must address in order to be subjectively rational. However, if such is the case, then it certainly has an impact on whether or not they count as knowing. This is precisely the sense in which luck does in fact play some role in a fallibilist account of knowledge.

Thus, the straight rule of actuality is fine for knowledge conditions. If it turns out that a skeptical scenario is in fact actual, then one cannot count as knowing. This conclusion would fit well with what we have said so far, because our grounds for treating skeptical hypotheses as irrelevant are that they are merely logically possible. However, if the skeptical hypothesis in fact holds then it is much more than logically possible; it is actual. We might say of agents who are completely unaware of a skeptical scenario that is actual that they are subjectively rational but that they fail to count as knowing.

The rule of resemblance needs to be modified in order to be consistent with a fallibilist view of knowledge. Resemblance is only relevant provided that the resembling scenario is sufficiently grounded. This is not as ad hoc as it might appear at first glance. If we recognize that we cannot entertain every single alternative and pursue truth sufficiently, then we must have some guidelines for treating alternatives that are not

sufficiently grounded as irrelevant whether or not the agents in the scenario can produce some reasons for treating them as irrelevant. To ask them to produce reasons is to ask them to treat the scenario as relevant, which has the effect of derailing the pursuit of truth. The argument given in the previous chapter is meant to demonstrate the extent to which the skeptic is asking us to derail our pursuit of truth in order to avoid error.

5.5: Epistemic Perspectives and Closure:

The solution I have pursued in the last two chapters raises some issues regarding Epistemic Closure. Relevant alternative views generally involve re-characterizing closure. However, as with many other relevant alternative views we will be able to save some version of the closure principle. If knowledge ascriptions are made in the context of a set of relevant alternatives, then closure will fail in instances in which the set of relevant alternatives shifts between premises. Alternately, we might claim that closure holds provided that the set of relevant alternatives remains fixed through-out the argument. There are different kinds of cases regarding closure and closure seems to hold in some while not in others. The cases of interest in the context of skepticism, cases like Dretske's zebra example, are cases in which closure seems counter-intuitive. These cases are cases in which the entailment under question is a skeptical alternative, and I will argue that Dretske is right to think that closure fails on the grounds that the evidence that grounds p is not sufficient to ground belief in q .

Closure can be articulated as the following argument form:

Closure: S knows that p
 S know that p entails q
Therefore, S knows that q

I have challenged lines of reasoning like the above when they involve shifting from a third person line of reasoning into a first person context. It is important to note that in the above argument, the first premise is a *knowledge ascription*. I have argued that knowledge ascriptions have distinct criteria from the actual possession of knowledge. So whether or not closure holds depends largely in the way in which the first premise is made. The mistake in the McKinsey paradox, for example, is not necessarily in the form of the reasoning from semantic externalism to knowledge of the contingent aspects of the external world. The mistake is made if an actual person uses such a line of reasoning to assert the first premise.

Dretske grounds his anti-closure argument in the zebra example that serves as counter-example to the closure principle:

Something's being a zebra implies that it is not a mule... cleverly disguised by the zoo authorities to look like a zebra. Do you know that these animals are not mules cleverly disguised? If you are tempted to say "Yes" to this question, think a moment about what reasons you have, what evidence can produce in favor of this claim. The evidence you had for thinking them zebras has been effectively neutralized, since it does not count toward their not being mules cleverly disguised to look like zebras.⁷⁴

Peter Klein argued that Dretske's argument fails because he makes an unnecessary assumption regarding evidence and closure. Indeed, Dretske's argument against closure is grounded in his claim that our evidence for x being a zebra does not "count toward their not being mules cleverly disguised to look like zebras". Klein's argument is essentially that if this assumption is not necessary for closure, and it seems not to be in several cases, then the assumption is unmotivated and Dretske's argument against closure fails. Klein argues that Dretske's counter-example holds only if "the closure principle entails that the very same evidence that justifies S in believing that the

⁷⁴ Fred Dretske, *Epistemic Operators*, 138-139.

animals are zebras must justify S in believing that they are not cleverly disguised mules because, it is presumed, that it is the only evidence we can be sure S has".⁷⁵

Closure holds as a fine form of reasoning as long as the first premise is asserted in a third person way of some potential knower. If asserted in a third person perspective the reasoning clearly holds. However, there are concerns if the first premise is asserted of an actual knower in a first person way. The concern stems from the fact that knowledge ascriptions are asserted because a knower exhibits *the indicators* for knowledge possession. The indicators, however, do not guarantee the possession of knowledge. Thus, there can be instances in which using the above kind of argument may be epistemically problematic in the first person because it involves an assumption that the knowledge ascription implies knowledge when such does not appear to be the case.

For example, there may be tempting to give arguments like the following:

Example: P1: I know that X is a zebra
P2: If X is zebra, then X is not a cleverly disguised mule.
C1: Therefore, I know that X is not a cleverly disguised mule.
C2: Since I know that X is not a cleverly disguised mule, I am entitled to ignore any evidence to the contrary.

What is problematic here is not the externalist line of reasoning that allows us to conclude that p entails q and that if we know that p, then we know that q. The problem lies in the tentative first person knowledge ascription that I know that p. If my knowledge ascription itself is tentative, closure can be problematic if one is tempted to use it as described in the example above. In the above example, such an agent would be prone to dogmatism in the first person in a way that is problematic. The tentative nature of knowledge ascriptions requires that we continue to be sensitive to error. However, an

⁷⁵ Peter Klein, How A Pyrrhonian Skeptic Might Respond, 38.

overzealous application of deductive closure may incline one towards dogmatism in a way that violates important norms regarding belief formation.

So in this sense, we can say that deductive closure holds in a third person way, and holds in first person contexts as long as the agent's application of deductive closure does not result in the violation of important first person epistemic norms. Namely, if the evidence that grounds a knowledge claim of p is sufficient to ground a knowledge claim of q, then there seems to be no problem. If such is not the case, then closure amounts to a form of reasoning that allows the conclusion to be on significantly stronger epistemic grounds than the premises. Where does such increase in epistemic strength come from if not passed through the premises? The only grounds for the conclusion of any argument are to be found in the grounds for believing the premises. This is why Dretske is absolutely correct to identify a problem with closure in the Zebra example. The rest of the reasoning for the argument beyond the grounds for P1 only applies in the instance in which S knows that p. Therefore, the entire line of reasoning is ultimately grounded upon S's claim to know that p. It is therefore, absolutely pertinent to think that the only grounds S has for believing that q are the same grounds that S has for believing that p.

So in this light we may characterize the problem with the McKinsey paradox in the following way.⁷⁶ The first premise is a knowledge ascription. Knowledge ascriptions are made on the basis of indicators for knowledge possession that do *not imply* knowledge possession. However, the deductive nature of the closure argument requires the actual possession of knowledge. If closure is based on a knowledge ascription that does not entail knowledge, then we have opened the door to the possibility of true

⁷⁶ For a full discussion of the McKinsey paradox see Martin Davies, "Externalism and Armchair Knowledge," in *New Essays on the A Priori*, (Oxford: Oxford University Press, 2000) 384-414.

premises and a false conclusion; such is the nature of shifting the principle from one that is formal to one that is informal.

Thus, the McKinsey line of reasoning holds in a third person way, but is not something that any real knower can actually make use of. My sense is that the application of such a line of reasoning actually violates requirements regarding justification in the first person. To assert that we know contingent things about our world based on a knowledge ascription about the meaning of our words, combined with a tentative philosophical theory, does not amount to justification in the first person. In a philosophical sense it may be an interesting argument, but we have standards for justification regarding beliefs about the external world that such an argument will not meet.

We might put the concern in the following way: knowledge implies truth, but knowledge ascriptions do not. To claim that ascribing knowledge of *p* to *S* implies the truth of *p* is tantamount to claiming that we know with certainty that *S* knows that *p*. Such might be fine in instances of knowledge grounded in a priori reasoning. However, those are not the cases that concern us here. Here we are concerned with external world knowledge and claiming infallibility of knowledge and knowledge ascriptions seems highly contentious in this arena. In fact it seems that such a claim would be tantamount to the claim that our external world beliefs are grounded in such a way that we can rule out in a definitive sense *all* other alternatives, including skeptical hypotheses.

If knowledge ascriptions require an internal justification, then we cannot claim to know that *X* is not a cleverly disguised mule because we have neither reason nor evidence to ground such a claim. Dretske is right to ask us how we might defend such a

knowledge claim. The alternative would be to allow knowledge claims in cases in which an agent has neither reason nor evidence in favor of her belief. Making the case for such an approach regarding knowledge ascriptions would be difficult in light of the fact that such an approach conflicts with our practices regarding knowledge ascriptions. Our practices are generally such that claims to know not grounded in some way are thrown out. If someone claimed to know that p , and when asked them how they know that p , they reply “I don’t know”- it would be consistent with common practice to treat such a knowledge claim as false.

In general, the deductive nature of the closure principle will require that the conclusion drawn is somehow contained in the premises already given. In cases in which q is included in p already, even in cases involving external world matters, closure does not seem problematic. Consider a slightly modified version of one of Klein’s examples: If S knows that water is present, then S knows that a clear, odorless, liquid is present. Such a case is not problematic because the justification for knowing that water is present is sufficient to ground the knowledge claim that a clear, odorless, liquid is present. The entailment in this last case is contained in p itself. Thus, the evidence in favor of p will be sufficiently in favor of q as well. However, that x is not a cleverly disguised mule steps beyond what is claimed in “ x is a zebra”. That this latter point is so may well act as straightforward evidence that knowledge claims are made against a set of relevant alternatives that does not include each and every alternative possibility there might be.

It is clear that this is not a full treatment of this issue, but having said the above we might be able to conclude a couple of things about the view I have developed here and closure. Closure will hold on my view in a third person way and it will hold in

Dretske's sense of closure; it will hold when S's evidence for p is sufficient for q. Otherwise, closure seems to fail.

5.6: Other instances of mistakes in shifting perspectives:

If it is right to urge that we need an epistemological account of both perspectives, it will also require guidelines in navigating those perspectives. Such guidelines will be the sort of thing that agents when in the first person perspective will rely upon and use as a way of avoiding epistemic mistakes. A little reflection reveals that such mistakes are possible (and, again, I have argued that skepticism is just such a mistake). Both McKinsey's paradox and Putnam's anti-skepticism argument are possible examples of mistakes in shifting perspectives.

Putnam's argument that utterances of skeptical hypotheses, such as "I am a brain in a vat" are self-refuting is another example of moving from a third person line of reasoning to a first person justification in a way that is epistemically problematic. Putnam's argument runs along the lines that any actual brain in a vat cannot refer to brains and vats and therefore cannot truly utter that statement. In the brain in a vat's world the terms "brains" and "vats" will refer to whatever usually causes the internal sensory experience associated with brains and vats, but in a brain in the vat world it will not actually be brains and vats that typically cause utterances of "brains" and "vats". So his argument rejects the problem of skepticism on purely semantic grounds.

Let's take a more detailed look at a Putnam type argument. Ted Warfield defends a seemingly plausible version of anti-skepticism via semantic externalism.⁷⁷ Warfield characterizes his anti-skepticism position in the following argument:

- P1. I think that water is wet (or, I know that I think that water is wet).
- P2. No brain in a vat in an otherwise empty world can think that water is wet.
- C1. So, I am not a brain in a vat in an otherwise empty world.

So this argument relies on an application of the closure principle. Since thinking about water implies that one is not a brain in a vat, Warfield moves from his claim to know that he is thinking about water to the claim that he knows he is not a brain in a vat.

This kind of move serves as further example of either problematic instances of the application of closure or of problematic instances of shifting epistemic perspectives.

Notice that premise P1 is a knowledge ascription. However the closure principle underlying this argument requires the actual possession of knowledge, not a mere knowledge ascription. Let's look at it this way. Knowledge ascriptions, even of the contents of our thoughts will not *entail* knowledge possession, at least in instances in which the contents of our thoughts are determined by externalist conditions as characterized in semantic externalism. As a quick example of this we might imagine that a brain in a vat has water like thoughts and says to herself "I am thinking that water is wet", and yet, she is not thinking that water is wet. Thus, the ascription of knowledge will not imply knowledge. However, the closure principle requires knowledge in order to be plausible.

To see the above let's imagine the closure principle as it actually occurs in the above anti-skepticism argument. Closure is occurring above as follows:

⁷⁷ See Ted Warfield, "A Prior Knowledge of the World," in *Skepticism* (Oxford: Oxford University Press, 1999) 76-92.

Instance of Closure:

P1. I claim to know that p.

P2. p implies that q.

C3. Therefore, I claim to know that q. Or, in other words, my justification for claiming knowledge of p transfers to a claim to know that q.

This is closure as it is actually applied, or it is an example of an *instance* of closure. An instance of closure is actually based upon knowledge *ascriptions*. Knowledge ascriptions are generally made in the first person and based on knowledge indicators. However, as we have seen, knowledge indicators will not imply or entail knowledge. There is nothing wrong with the instance of closure argument. However, focusing on the fact that any instance of closure is grounded in a knowledge claim allows us to see that the knowledge claim gained at the conclusion is no more decisive than the justification for the other knowledge ascriptions made thru-out the argument.

Perhaps it has traditionally been thought that knowledge ascriptions of contents of thoughts do imply the possession of knowledge. We can see from the example above that such is not the case. However, this gets tricky with semantic externalism. It is difficult to capture how the brain in the vat is wrong about her thoughts on an externalist view of the contents of her thoughts. Thus, perhaps we should concede to the Putnam camps that even the brain in a vats' knowledge ascriptions do imply knowledge. On this view it becomes impossible to articulate or capture in the utterances of the brain in a vat how she is wrong about the nature of her environment. And, yet, it seems so clear that she is massively deceived and wrong. This massive deception is at the heart of skeptical scenarios. The brain in a vat thinks she has a body, and she does not, and she thinks there are trees, lakes, grass, and there is not. Yet, it is impossible to characterize this epistemic error in utterances or in the thought contents of the brain in a vat on an externalist

account. Yet, it seems that even on the externalist account the epistemic problem of skepticism remains. It is undeniable that the brain in a vat suffers from a state of extreme epistemic poverty about her environment. So, admittedly it becomes somewhat difficult to characterize the skeptical problem in the context of semantic externalism. It is not quite correct to say that the brain in a vat thinks there are trees, lakes, hands, and she is wrong about such things.

While I do not want to lobby major criticism in a quick or off-handed way we might make some notes about the implications of the above view. If it is true that the brain in a vat's beliefs are in fact true, such a view may well imply that brains in vats do in fact possess knowledge of their external world.

Yet it seems counter-intuitive to ascribe external world knowledge to beings who could not accurately describe their external world. Brains in vats have no real understanding of the nature of their world. It just so happens that their beliefs are true because of how semantic externalism ascribes content. But the truth of the brain in a vat's beliefs seems accidental in exactly the way that is inconsistent with knowledge. In fact these implications seem to embody the worst criticism against externalist positions in general because it is not possible to capture how deprived the BIV is despite the fact that she has no real understanding of her environment. Such worries are the embodiment of the concern that externalism allows knowledge even when the internal state of the agent is epistemically lacking.

Notice even further, that semantic externalism may well imply that a brain in a vat's sensory experience is reliable. Even in skeptical scenarios, sensory experience consistently produces true beliefs because on many such views content is assigned

according to whatever typically causes such utterances. Thus, sensory experience is reliable in the sense that it regularly produces true beliefs. There becomes this strong sense in which the semantic externalist is actually denying the *metaphysical* possibility of massive deception. Massive deception is impossible because *whatever* triggers our utterances will determine content, thus it seems impossible to be wrong in the sense of having a mismatch between content or thoughts and the causes of the thoughts. However, such a mismatch is exactly the possibility that skeptical scenarios aim to capture. If these concerns do not amount to outright reductios of this brand of semantic externalism, then such implications at least show strong prima facie concerns about the implications of semantic externalism in epistemology. Such concerns might be so strong that they must be worked out for the theory to be viable.

So in light of this dilemma we might try a different look at the Putnam type argument. Notice there is a slight circularity built into Warfield's argument. His first premise is a knowledge claim about thinking about water, which implies that he is not a brain in a vat. However, notice that not being a brain in a vat is a necessary precondition for the truth of the first premise. In order to think about water, he cannot be a brain in a vat. Thus, he derives one of the necessary preconditions of his first premise as conclusion. This works only if the truth of the first premise is so apparent as that it does itself imply the conclusion. But what are the grounds of the first premise? Having water like thoughts? Surely it is not impossible for the brain in a vat to have water like thoughts. Is it transparent from water thoughts which environment one is actually in?

To my mind the crux of the problem here is the shifting of epistemic perspectives without any sensitivity to the nature of those perspectives. The philosopher here moves

from a highly externalist line of reasoning about the conditions that determine the contents of thoughts and then combines that with an assumption that all of those conditions are transparent from the first person perspective. It seems that if we are going to be externalist about the contents of thoughts we may well have to accept that at least some of the conditions that determine the contents of our thoughts are beyond the perspective of actual thoughts. Indeed, such external conditions on content are very similar to external conditions on knowledge. If such is true, then closure may well be a valid argument in the third person, but potentially problematic in the first person.

Closure in the first person becomes a transfer of justification or a transfer of warrant. The result is that the knowledge claim at the end of an application of closure is justified if the warrant that transferred from the first knowledge claim is sufficient justification for the consequent. So if p implies q , this will not matter if the warrant that justified the assertion of knowledge for p does not justify the assertion of knowledge for q . So in many instances closure will be acceptable because the justification for p serves as sufficient justification for q as well. However, in other instances the justification one has for p will not serve as the kind of justification that supports a knowledge claim of q .

It seems to my mind that it is remarkably important that we keep in mind that knowledge ascriptions function differently than knowledge. As we have mentioned, knowledge may not require an internalist type justification, but, on the other hand, it seems highly plausible that knowledge ascriptions do. If one claims knowledge, surely one must have a reason to think she in fact possesses knowledge, and such a reason should be accessible to the agent via introspection.

5.7: Conclusion:

So we have reached a point where I have conceded to the skeptic that we function on the assumption of the reliability of the senses in our first person endeavors. The skeptic, as we know, would claim that such a justification is not warranted. Since our assumption is unwarranted, beliefs justified in a way that relies on this assumption fail to amount to knowledge. It was argued in the previous chapter that our assumptions of reliability regarding sensory experience are justified for the sake of pursuing truth. I have framed the skeptical argument as a reflection of the value of epistemic goals, and I have argued that the goal of pursuing truth outweighs the value of avoiding error. This argument is aimed at undermining the skeptic's central claim that agents who rely on sensory experience are somehow epistemically blameworthy or deficient. The classic skeptical charge is that agents who rely on sensory experience are somehow epistemically naïve. We have seen, however, that skepticism can be construed as a choice. The choice is between the pursuit of truth with the possibility that we will get it wrong, or to avoid all possibility of error and forsake the pursuit of truth. There really is no way to avoid this choice. One either relies on and endorses the information given to us via sensory experience, or not.

Granted, there may be fine grained distinctions about how full bodied the reliance on sensory experience should be, but I have argued that such fine grain distinctions are merely splitting hairs. The epistemic debate comes at the level of whether or not we have justification to assume the reliability of sensory experience in the first place, and it has been argued that we do in fact possess such justification.

However, our justification for relying on sensory experience is not classic epistemic justification; our justification does not count toward the truth of our assumption in any way. I have defended position on the grounds that what justifies our reliance on sensory experience is a meta-principle regarding the pursuit of truth in the context of fallibility. This meta-principle allows us to endorse methods of belief formation for the sake of pursuing truth. The AR meta-principle is looser than the skeptic's requirement; the AR principle says we are justified in relying on methods of belief formation until there are indicators of unreliability. This meta-principle maximizes the attainment of truth while minimizing the most pernicious and likely error. However, we must admit contrary to the skeptic that our principle and our stance does not provide certainty; instead, it allows us to pursue truth despite fallibility. Thus, a boiled down version of the defense of AR is that our reliance on sensory experience is grounded in the value of and pursuit of truth.

This justification is not pragmatic in the traditional sense. The traditional pragmatic versus epistemic division has been made between that which involves the pursuit of truth and that which does not. Epistemic justification has traditionally been that kind of justification which counts towards truth. And pragmatic justifications have generally been those regarding our non-epistemic goals. Thus, our justification is not clearly epistemic or pragmatic in the traditional sense of this distinction. However, I have argued that since AR results in the attainment of more truth that it should count as an epistemic justification. Does the justification for relying on sensory experience provided indicate that our sensory experience is in fact reliable? No, clearly it does not. Rather, it indicates that a reliance on sensory experience results in more truth. The nature

of this justification seems acceptable in light of the fact that the skeptic's approach excludes the attainment of external world truths for all agents, even those who could attain truth if they followed AR. This is the sense in which knowledge and the pursuit of truth does involve some element of luck after all.

Cartesian Skepticism as Moral Dilemma

“Remaining Questions and Concluding Remarks”

Chapter 6.1: Summary:

I have argued that in many everyday contexts we have good reasons for ignoring skeptical hypotheses. The simplest version of the argument I have given is that if we are going to attain truths about the environment in which we live, we can only do so by relying on sensory experience. Thus, the pursuit of truth requires us to accept sensory experience as reliable. Granted, we need not, and should not, have blind faith in sensory experience, or any method of belief formation for that matter. However, we cannot sit back and wait for the proof the skeptic seeks without severe epistemic consequences. The skeptic's position involves the loss of truth even for those agents for whom external world truths are possible. So the pursuit of truth itself involves some possibility of error, but this should not surprise us. We have long been aware of the fact that we are highly fallible knowers.

So we may sum up the argument in the following way. First, we have focused along with Bonjour on skepticism as it relates to the first person perspective; as it relates primarily to applied knowledge ascriptions. Secondly, I have argued that first person practices regarding relevant alternatives are epistemically justified rather than merely pragmatically justified. This conclusion allows us to dismiss mere logical possibilities in the context of everyday life. This will include dismissing skeptical hypotheses.

The crux of this argument has been that our first person practices are largely goal and value oriented. One of the primary values is truth. This creates strong first person norms governing the attainment of truth and avoidance of error. The result is that we must balance these competing concerns in the first person. If we are overly cautious we are likely to sacrifice many truths. If Descartes had not found a way to come out of the skeptical hole he dug for himself, he would have missed out on much knowledge and many truths. Clearly, such an example is not epistemically ideal. Thus, our leap of faith is really an epistemic one, and one that is not merely the result of undeniable psychological or practical forces.

Admittedly, my argument rests on a number of claims that should be made explicit. First, I am relying on an analogy between moral values and epistemic ones. In ethics the well worn example of the refugees hiding in the attic and the Nazis knocking at the door illustrates that moral obligations can conflict in such a way that any action we choose will land us in violation of a moral obligation in the context. In the aforementioned example, agents in such circumstances must choose between telling the truth and saving innocent lives. In much the same way, I have claimed, epistemic values conflict. The circumstance of this conflict forces a choice that aligns either with the pursuit of truth or the avoidance of error. Granted, perhaps there are very subtle and nuanced positions in the middle, but, as argued in the Introduction and elsewhere, these nuanced positions do little to avoid this choice. The nuances and subtleties to be found that may be lacking in this dichotomous choice embody subtly that does not avoid the basic tension and conflict as I have characterized it thru out our discussion.

It seems safe to conclude that if we are discussing skeptical hypotheses seriously we are in a philosophical context. I have argued that the epistemically important distinction between philosophical contexts and everyday life has to do with the difference in epistemic perspectives between these two contexts. Since there are important epistemic differences between philosophical and everyday contexts, we have some non-question begging grounds for taking skepticism seriously in a philosophical context but dismissing it in the context of everyday life.

To my mind, this conclusion fits our intuitions regarding skepticism very well. There may be times in which we seek as in-depth an inquiry as possible. In such a context alternatives are raised more easily than in other contexts in which depth is not the primary focus. Everyday contexts are such contexts in which the standards for raising an alternative are stricter. This is appropriate given the fact that the contemplation of an alternative takes time and energy. Everyday contexts are often contexts in which depth conflicts with the pursuit of truth in the time given; thought experiment G in chapter 3 shows what such a context might look like.

Lastly, this view allows us a stronger relevant alternatives position than previously thought. It is not as though we are ignoring skeptical hypotheses simply on the grounds that no one has mentioned one, nor is it the case that we are ignoring skeptical hypotheses for ad hoc reasons. It is not as though we make an exception to the Rule of Exclusion for skeptical hypotheses on arbitrary grounds. Our reason for dismissing a skeptical hypothesis is at the very heart of all our epistemic endeavors; our reason is for the pursuit of truth. As we saw in the Pascal Wager argument, it is not as though brains in vats are taking a leap of faith at the expense of potential truth. Unless

the epistemic situation of the brain in a vat changes, she is not capable of gaining truth about the contingent aspects of her environment. Thus, it is not as though being cautious to avoid error when regarding skeptical hypotheses gives us an avenue to truth. Thus, such alternatives are properly ignored for the sake of pursuing and possibly attaining truth.

6.2: Skepticism in a Philosophical Context:

Because we are often in a first person perspective in our everyday lives, and that in this perspective we have to balance the goal of truth with expediency and efficiency, we are often justified in rejecting skeptical hypotheses. By arguing such we have some epistemic reasons for dismissing skeptical hypotheses in most everyday contexts. Further support for this position is offered by the Pascal wager type argument we evaluated in the last chapter.

The question of skepticism in a philosophical context remains to be answered. I have largely conflated the third person perspective with a philosophical one. While this is not crucial to my argument, I have done this because a third person perspective is concerned with such uniquely philosophical questions as: what is it to possess knowledge, what is constitutive of justification, when are knowledge ascriptions appropriate, etc. So let us briefly look at how we might respond to skepticism in a philosophical context in light of our argument thus far.

The first thing to note is that the response to skepticism developed thus far will not pertain to issues of skepticism in a philosophical context. As the dialectic has been

construed, the solution offered here is effective for the issue of knowledge ascriptions, since they are primarily carried out in the first person. Thus, skepticism in a philosophical context concerns questions about how knowledge is possible in general, how it is possible in light of the fallibility of sensory experience, and how knowledge is possible in light of skeptical hypotheses, etc.

While it is not my aim to offer an answer to these questions, there are two things that are worth mentioning. First, a response to philosophical skepticism should occur in the context of a full epistemology. Pritchard's externalist response serves as a good example of ways in which we might handle philosophical skepticism. My own sense is that externalist type responses do well against skepticism precisely because they require less in terms of transparency of evidence in the first person. As previously discussed though, such a view has correctly been seen as failing to answer the first person issue of whether or not *our actual* ordinary beliefs about the world *are in fact* justified. So it has been this issue that I have focused on in our previous discussion.

The second thing that I want to mention about this issue is that there is a large sense in which skepticism raised in a philosophical context poses significant challenges to anyone articulating a theory of knowledge in which knowledge is a regular occurrence. To my mind, this is an appropriate stance to have toward skepticism. The mistake lies not in thinking that skepticism has immense philosophical interest, but in assuming that the philosophical difficulties translate into our everyday lives in a way that is not epistemically problematic. Philosophers, perhaps not surprisingly, have been insensitive to the important epistemic differences between philosophical contexts and everyday life.

In this latter sense, the argument given here does address some of this philosophical skepticism. The argument under consideration here is one that gives some idea of how knowledge might be seen as fallible in ways that are not overly counter-intuitive. As discussed previously, the view articulated here need not be a brand of highly counter-intuitive fallibilism. The fallibilism considered is one in which agents claim knowledge with the belief that there are no other serious and relevant error possibilities that are un-eliminated.

In his “*Epistemological Realism*” Michael Williams characterizes the ways in which we might be dissatisfied with possible responses to skepticism. Even if we avoid skepticism in the first order, he notes, we find various ways in which we might continue to be dissatisfied: “We may *have* knowledge of the world, but we will never be able to explain to ourselves *how* we do. We may know things about the world, but we will never know that we know them.”⁷⁸ How does the solution proposed here stack up against these concerns? In a straightforward sense I think we can *see how* knowledge of the external world is possible. The second question is more pertinent for the response I have suggested. Without endorsing the standards he seems to evoke, it seems he thinks of knowing that we know as having some conclusive reason to think we count as knowing. However, if we are fallibilist about knowledge, we may well want to be fallibilists about second order knowledge as well. I have endorsed a view in which we are justified in claiming that we know because we possess reasons to think we have filled the first order criteria necessary for knowledge possession. Yet, I have claimed such despite the fact that some of the criteria at the first order level, some of the criteria we must fill in order to possess knowledge, is essentially externalist in nature. Specifically,

⁷⁸ Michael Williams, “Epistemological Realism,” 54.

if we are to count as knowing it must turn out that our senses are in fact reliable; that we are not brains in vats.

Whether or not we want to endorse Williams' more stringent standards, we might ask whether or not the solution I have suggested meets this standard, and we must note that it does not. In the sense Williams seems to endorse, knowing that we possess knowledge of the external would seem to require knowing that we are not brains in vats. However, despite our endorsement of AR we do not know that we are not brains in vats. In fact, I have argued that we are entitled to endorse AR as a meta-principle, and as such we can endorse it despite that we fail to know which possible world is actual. Since I have argued that we are warranted in asserting knowledge, since we seem to have met the criteria for knowledge possession, it will follow that on the view I have put forward, warranted assertability is not the same as knowing that you know as Williams seems to characterize it.

6.3: Moore and Skepticism:

I have largely characterized the argument I have given as a relevant alternatives argument. However, it also seems proper to treat our defense as in the spirit of the Moorean response to skepticism. Moore's common sense rejection of skeptical hypotheses is exactly the attitude of the everyday knower that I have sought to defend as epistemically sound.

Moore's grounds for dismissing skepticism lie in how poorly skeptical hypotheses and arguments are supported. Moore questions the skeptic's contention that we ought to

abandon that which we are sure of (such as having hands) in favor of that which we are not (fancy philosophical argument).

In “Proof of an External World” Moore notes that skeptics do not seek proof of statements such as “I have hands”, but, instead, seek a more general proof of how such statements can be known at all, or proven at all.⁷⁹ Moore admits on this score that “I have conclusive evidence that I am awake: but that is a very different thing from being able to prove it.” Much of what has been argued here is grounded in a notion that we have evidence in favor of our sensory experience, but not the kind of evidence that speaks against skeptical alternatives.

At this point it might be worth noting why it might be that the skeptic does not count any of the evidence that ordinary believers would count in favor of the reliability of sensory experience. Many of us have experiences involving the confirmation of our sensory experience; it is often confirmed by those around us, when relied upon its’ veracity is often demonstrable, etc. Yet, we must admit, as Moore had to admit, that this kind of evidence is not of the sort the skeptic seeks. However, there is a bit of a slight of hand on the part of the skeptic. The skeptic operates with a standard of certainty embodied in the requirement of ruling out the various skeptical hypotheses.

As we have previously discussed, skeptical hypotheses are merely possible error. The skeptic has no evidence to ground the claim that we are dreaming, that we are BIVs, or that sensory experience is *systematically* unreliable. Thus, skeptical alternatives do not express the worry of likely error; they express the worry of possible error. In the

⁷⁹ See Moore, “Proof of An External World,” in *Epistemology*, ed. Ernest Sosa et al, (Malden: Blackwell Publishing, 2008) 26-28.

skeptic's mind knowledge involves certainty, how can one know if one is not certain the skeptic wonders?

However, much like the distinction between the concept of formal validity and that of logical strength, if the skeptic's standard is certainty and we fail to meet that requirement, then according to her we have *no grounds* for preferring our ordinary beliefs or for relying on our sensory experience. Yet, this is a bit fast. The formal concept of validity is a black and white concept that involves a precision stemming from the standard of certainty in formal logic. Thus, if an argument does not provide certainty, then in formal logic it is invalid. The concept and the standard it embodies does not admit of degree.

If we reason with the skeptic then we have to admit that we have *no evidence* and *no epistemic grounds* for preferring our ordinary beliefs about the world. However, if we do not embrace the standard of certainty and the black and white concept of validity that it entails, we can admit of degree. We can admit that we have much evidence and good reasons to rely on sensory experience, but they are grounds that do not provide certainty. Thus, it can be that on the skeptic's view of knowledge, we have no grounds for preferring our ordinary beliefs, but on a fallibilist view of knowledge we indeed have grounds that provide some likelihood of truth. Whatever those grounds may be, they will not compel the skeptic. The skeptic gets her conclusion that we have no grounds for preferring our ordinary beliefs because we have no evidence against skeptical hypotheses and in favor of our ordinary beliefs, but this requirement embodies a standard of certainty and it is only on that standard that we fail to possess any reasons to ground our ordinary beliefs.

The above is the sense in which the approach I have taken towards skepticism cannot provide a refutation of skepticism anymore than Moore's arguments provided a refutation of skepticism. And, yet, there is something undeniably appealing in Moore's response. One has to work themselves into a skeptical philosophical context in order to find the skeptic's argument compelling; this seems true even for the skeptic. I have wanted to do justice to the de facto practice of relying on our sensory experience that we all engage in and find impossible to deny. If we find it impossible to do justice to these practices in a philosophical way, in a way that engages the skeptic's demands, such may be as much an indication of a deficiency in theoretical epistemology as it is in the practices everyday agents engage in.

6.4: Sensory Experience, Knowledge, and Luck:

We might wonder at this point, though, whether or not I have really defended the everyday attitude towards sensory experience. I have presented the endorsement of sensory experience, particularly so in the context of the Pascal's Wager, as though it occurs in the spirit of a bet. However, this cannot possibly be right. We believe that sensory experience is reliable in a much stronger way than as a tentative and rational bet given our circumstances. I have to admit such criticism would be well founded. It is not as though we are tentative in our reliance on sensory experience. However, we have evaluated the assumption of the reliability of sensory experience in a much more detailed and philosophical light than is common in everyday life. Additionally, as talked about earlier, we have approached relying on sensory experience as an epistemic issue only, when, clearly, our reliance on sensory experience extends beyond the epistemic.

It has been convincingly argued that belief grounded in sensory experience is not the product of some rational process reflecting conscious endorsement of the given content. Beliefs grounded in sensory experience, seem, as a matter of fact, to arise spontaneously with experience. However, the skeptic has asked us to step back from this process and determine whether or not this endorsement of sensory experience is in fact justified. The skeptic has argued that without some reason to rely on sensory experience we have no reason to prefer ordinary beliefs over skeptical rivals. Thus, we too have stepped back from whatever process may in fact take place to evaluate said process from an epistemic point of view. Having done such, I have argued that despite that we are in some sense “blind” we are justified in relying on sensory experience. We have granted as much as might be reasonable to grant the skeptic and found that even in light of such concession we still have reason to prefer relying on sensory experience over doubting it.

6.5: Contextualism, Relevant Alternatives, and the Rule of Exclusion:

I have branded the argument I have given as a sort of Relevant Alternatives position. The argument given has the features of a relevant alternatives argument. I have treated skepticism as an argument stemming from the fact of skeptical alternatives, and the skeptic’s claim that we fail to possess sufficient evidence to rule them out. I have argued, on the one hand, that it is the concrete contextual features of first person belief formation that rule out skeptical alternatives as relevant contenders. Alternately, I have argued beyond the above brute-fact to the conclusion that it is for compelling epistemic

reasons that skeptical hypotheses should not be treated as relevant alternatives in most everyday contexts.

Despite the contextualist nature of the argument given, it should be noted that my argument is not grounded in a claim that the standards for knowledge change. Nor is it the case that my argument is grounded in some implicit claim that there are weak and strong kinds of knowledge. Additionally, I do not think my argument need be construed as one in which the standards for evidence shift according to context.

It might be easy to read my argument as one in which it happens that in everyday life our standards for evidence are low because of the concrete aspects of belief formation in the context of everyday life. Alternately then, it would be easy to think that in the philosophy classroom we raise the standards for evidence and produce some stronger version of knowledge grounded in such reflection. I think such a reading mistaken and unnecessary. The consequent of such an interpretation is that the product of such deep reflection stands on *significantly* stronger epistemic ground. This consequent stands in conflict with the primary thesis of my argument- that everyday knowers are on perfectly sound epistemic ground. Instead of attaining some kind of *strong* knowledge after philosophical reflection on our beliefs, we attain some kind of relief and comfort in the knowledge that our practices regarding sensory experience and our usual dismissal of skepticism is indeed justified.

Let me offer an interpretation that avoids the above implications. First, the standards for knowledge are not shifting on my view because they remain consistent across an everyday and philosophical context. I have agreed with the skeptic that in order to claim to know our ordinary beliefs, we must have reason for ruling out the skeptical

alternatives. I have disagreed with the skeptic in her implicit claim that the ruling out of skeptical hypotheses must occur on the basis of evidence that speaks against skeptical hypotheses. Instead, I argued that the concrete aspects of the first person are such that alternatives generally have to be on stronger epistemic ground than merely logically possible.

It should be noted that I do not want to ground my position in the claim that *standards for evidence* shift when we move from an everyday to a philosophical context. Admittedly, the concrete parameters of belief formation, such as time constraints, shift when we move contexts, and, admittedly, we shift perspectives and focus on more general questions surrounding knowledge and justification when we move into philosophical contexts. However, this should not be construed as a shift in the standards for evidence. Rather, it seems what actually occurs is that we must make decisions and form beliefs under time constraints in everyday life, and while we make what we take to be well founded decisions, we occasionally long for a more thorough and in depth evaluation of our justifications, beliefs, and reasoning processes. Thus, what I see as different between everyday life and a philosophical setting is that we remove the concrete parameters of everyday life and shift our focus to more general questions. We ask the questions we often do not have time for, and we review decisions made quickly.

A philosophical context allows us to evaluate in detail whether or not we are justified in neglecting skeptical hypotheses the way that we do. It allows us to open a door we had closed in order to determine whether it really is best that it stay shut. What we come to at the end of such reflection on the justification we possess for our ordinary beliefs is that we are in fact well grounded in our epistemic practices, they are both

rational from a subjective view point and justified from the objective requirement of truth-conducive. If after philosophical reflection we come to doubt that we are justified in our ordinary beliefs it is only because we have consented with the skeptic that not only do we have to rule out skeptical hypotheses, we must do so *in light of* evidence that speaks directly against them (skeptical hypotheses) and in favor of our ordinary beliefs.

The only sympathy I can garner for such a view is founded in our usual practice regarding the Rule of Exclusion. Since contending alternatives are usually ones that are on stronger grounds than mere possibility, it is usually the case that we require, or strongly prefer, *evidence* that rules out contenders. Given that skeptical hypotheses themselves are on weak ground it seems difficult to motivate such a strong requirement in order to rule them out. They are merely logically possible, and as such it seems perfectly reasonable to rule them out on concerns about attaining possible truths. It seems fine if we are willing to allow skeptical alternatives that we rule them out on the basis of a cost/benefit analysis. If we know that relying on sensory experience is an epistemically sound practice regardless of which world we are in, then our ordinary beliefs grounded in sensory experience are well founded. That there is a possibility of error only shows that knowledge is fallible and that there is some element of luck involved in epistemic endeavors after all. Knowledge is not attained solely on the grounds of the subjective efforts of the agent herself. No matter her efforts, if the tools she has to work with are, unbeknownst to her, faulty, then she cannot attain knowledge. Skepticism is worrying or upsetting only because it articulates a scenario in which something of tremendous intrinsic value, namely truth, is inaccessible. It is upsetting because deception is always an upsetting prospect.

6.6: Other Solutions to Skepticism:

The project here has been to do justice to the intuition that the everyday knower is motivated by deep epistemic concerns rather than mere psychological ones in a way that allows us a *philosophical* defense of our first person practices. While there are hints of this kind of response thru out the literature on skepticism, there has been very little serious endeavor to articulate how our everyday epistemic practices are in fact justified in an *epistemic* way, rather than in a merely pragmatic way.

There are advantages the approach articulated here has over its competitors. Let's first think about the basic contextualist view of the sort Keith DeRose describes as allowing for a powerful attack of the skeptic, while still allowing for the persuasiveness of the skeptical argument.⁸⁰ Such a contextualist view, in simple outline, maintains that the standards to count as knowing shift from context to context. On such a variantist position it will turn out that in some contexts the standards to count as knowing are relatively low while the standards are much higher in other contexts. Such a view thus allows the contextualist to maintain that we count as knowing in everyday contexts in which the standards are relatively low but fail to count as knowing in skeptical contexts in which the standards have been raised significantly.

First, let me note what seems to be right in the basic view outlined above. The contextualist seems right about the fact that the standards employed by those in everyday life are different than those the skeptic employs. The skeptic's demand that any agent claiming to know must handle all and any error scenarios seems tantamount to requiring a standard of certainty in order to count as knowing. It also seems that we generally do not employ such high standards in everyday life. However, what the contextualist sees as a

⁸⁰ See DeRose, "Introduction," 17-18.

mere difference in context strikes me as more blatant disagreement between the skeptic and the everyday agent about the nature of knowledge.

Notice that in a theoretical sense the contextualist has in somewhat skirted the skeptical problem by assuming that all various standards are acceptable.⁸¹ The skeptic would maintain that the lower standards of everyday life are not examples of true knowledge and that we are wrong to count them as such. To this criticism the contextualist has no immediate response. The skeptic would simply note that in order to count as knowing an agent must possess some evidence against skeptical hypotheses but she has none whatsoever. The general contextualist response has done little to motivate the infallibilist (and variantist) view of knowledge being endorsed in the contextualist solution. While some have tried to motivate the variantist aspect of knowledge in contextualism via arguments about indexicals, there has been little real work to motivate the infallibilist aspect of knowledge. Also, it is the infallibilist view of knowledge in contextualism that does the much of the anti-skeptical work. I have made some effort here to motivate an infallibilist view of knowledge rather than merely embracing it as brute fact. Notice that the contextualist has a largely concessive response to the skeptic. She has to admit that in skeptical contexts we have no response other than to concede to the skeptic. There is something counter-intuitive about making such a concession that we do not in fact know that we have hands, for example, while maintaining a short while later in another context that we do in fact know that we have hands. It seems the blatant contradiction here ought to be unsettling particularly considering that our evidence has not changed from one context to the next.

⁸¹ Of course, though, this is too strong and not something the contextualist is literally committed to.

The motivational argument provided here lies in acknowledging the epistemic value of what we seek in our epistemic endeavors and the recognition that the skeptic's high standards do come at an epistemic price. I have tried to motivate the view that in the face of fallibility we still have ample motivation to seek truth, and doing such requires that we jump in with both feet. I have suggested that the assumptions of reliability that lie underneath our pursuit of truth are rational assumptions, and that such assumptions are conducive to the pursuit of truth.

Notice too that I have picked up where Duncan Pritchard's work leaves off. He leaves us in his *Epistemic Luck* with a level of uncertainty regarding first person endeavors and skepticism. He argues that our first person endeavors are filled with a certain amount of epistemic *angst*; he likens us to tightrope walkers who assume but do not know that the safety net has been erected below. This description inevitably leaves us wondering why we are justified in assuming the safety net is there in the first place. What justifies this assumption? There are many solutions to skepticism, and many interesting solutions at that, but very few of them directly address the question about why we are justified in assuming sensory experience is reliable. The typical skeptical explanation would be that we assume sensory experience is reliable because we are either so psychologically compelled or we are irrational and unenlightened. The argument I have given here shows that such is not necessarily the case.

The assumption that sensory experience is reliable is both epistemically efficient and epistemically warranted. Such assumptions make sense in light of the fact that we do not know which world we inhabit and that we must, in a simplistic way, choose between caution to avoid error and recklessness to pursue truth. While I have encouraged the AR

position that strikes a balance between our two competing concerns, that of avoiding error and that of attaining truth, the AR position cannot entirely mediate the fact that in everyday life the pursuit of one goal has to come at a cost to the other.

The above epistemic reality should not surprise us too much. It is often the case in morality that the promotion of one moral good comes at a cost to another good. The pursuit of economic growth based on consumption, for example, is a good of a sort that may inherently involve a cost to the environment, a good of another sort. However, such is the nature of real life; it is messy and the parameters of real decision making force us to make choices we might not recognize theoretically. Thus, it might seem in a philosophical context that the avoidance of error is being recklessly ignored by the naïve everyday knower, when in reality the situation is more complicated than that. The everyday knower, aware of it or not, is in a circumstance in which epistemic goods cannot fully and simultaneously be pursued as we might wish to be the case in a theoretical context. Thus, one of strengths of the discussion here is that it aims to address these first person skeptical worries in a way that openly addresses the question of what justification we have for assuming sensory experience is reliable, and why we might want in a theoretical epistemology to embrace infallible knowledge.

6.7: Concluding Remarks:

My project has been in line with Bonjour's claim that at the heart of skepticism is a concern or worry that is generated about first person knowledge ascriptions. I have concurred with this assessment on the grounds that third person questions raised by skepticism are largely handled well by an externalist account of knowledge. However, in

concurrence with philosophers as different as Bonjour and Pritchard, such an approach to skepticism does little to handle those pertinent first person issues. However, unlike Bonjour and others, it seems that if we are concerned with our everyday beliefs and our everyday knowledge ascriptions, it is very important to approach the defense of these practices in a way that is consistent with actual everyday knowledge claims and actual everyday knowers. Fancy philosophical arguments do not defend the actual practices of everyday knowers who do not in fact defend their ordinary beliefs with such arguments.

The question is whether or not the practices we engage in a non-philosophical context are truly epistemically justified. My answer to this has been yes. And I have attempted an argument that is perfectly consistent with the practices of actual knowers, and aims to justify their practices without requiring such everyday knowers to appeal to any of the arguments contained within this paper.

Admittedly my goal has been to do justice to two strong and prevailing intuitions regarding skepticism. The first intuition is that the many people who respond to skepticism with annoyance, the many that are put off by such off-the-wall-hypotheses as brains in vats, are right. The second intuition is that it is a deep epistemological issue to understand how knowledge is possible in light of skeptical hypotheses. We are too quick to assume that our philosophical difficulties with skepticism translate into our everyday lives. This is the very concern that Bonjour raises regarding skepticism. His approach to skepticism centers on concerns over the status of our actual, everyday, ordinary beliefs about the world. I have argued that our everyday lives are governed by different epistemic norms than philosophical contexts are.⁸² There is a large body of evidence that

⁸² This is probably too quick. As clarification, the norms in the first person are different because of concrete aspects of belief formation in everyday life. Factors such as time, efficiency, and import in terms

indicates that knowledge is possible in everyday life (i.e. our many claims of knowledge, etc). And there is a large body of evidence that indicates that the contemplation of skeptical hypotheses is not always epistemically appropriate.

The approach I have taken to skepticism is in the spirit of Moore's response to skepticism. Instead of rejecting skeptical *argument* on the grounds of a lack of evidence, the position given here rejects skeptical *alternatives* on the grounds of a lack of evidence. The standard of taking seriously and ruling out, where ruling out is equivalent to knowing that the skeptical hypothesis does not hold, each and every possibility of error is tantamount to denying the formal and informal distinctions in reasoning. Descartes wanted the precision of formal reasoning and the power of informal reasoning; yet, we have seen through skepticism that such a combination is not possible. The power of informal reasoning comes at the price of inherent fallibility. Skeptical alternatives embody the limits of this kind of reasoning. Such a possibility of error cannot be refuted even with the best kind of evidence we might have for the knowledge in question. Yet this is not as counter-intuitive as some think. Claiming knowledge of p while admitting that it is possible, but not likely, that one might be wrong does not seem problematic. It would be a very different matter if skeptical hypotheses were grounded in evidence in one way or another. It seems much more counter-intuitive to claim knowledge while admitting a possibility of error that is likely.

Unless Epistemologists wish to undermine the very subject of our study, it is important that we develop a concept of knowledge that fits all that which is properly informal. I have argued that the skeptic has no *prima facie* grounds for ruling out

of consequences all affect the resulting epistemic norms. However, this does not mean that the epistemic *values* are any different. The primary epistemic values in both contexts are the attainment of truth.

fallible knowledge, and we have no clear reasons to reject the possibility of such knowledge. Rather, we seem to have much evidence that many, if not most, endorse and embrace a fallibilist position regarding knowledge in many epistemic arenas.

The skeptic relies on the Rule of Exclusion, and our psychological attachment to sensory experience to make her skeptical alternatives seem more epistemically powerful than they really are. Philosophers often talk as though many find an appealing argument in skepticism. However, this does not really seem to be the case. Most students in a philosophy classroom do not worry about skepticism after hearing the skeptic's argument; they dismiss it as far-fetched and often have the attitude that it is ridiculous to demand that we rule out such hypotheses. An inability on the part of Philosophers to clarify the theoretical concept of knowledge in such a way as that it and knowledge ascriptions are consistent with the fallibility articulated in skeptical hypotheses is not necessarily a sign of a deficiency on the part of everyday knowers.

Rather, the above failure on the part of epistemologists might be an indication of a tradition that has been overly influenced by standards of certainty when it comes to knowledge. Thus, while it may not be possible to knock down the skeptic when she is playing "King of the Mountain", we have offered some reasons for thinking that her long held stance as "King of the Mountain" is not warranted.⁸³ So while we might admit to the Descartes of the first meditation that we do not know for sure that we are not currently dreaming, we will also maintain that we do not have sufficient reason to worry that we are in fact dreaming to warrant even a cease- fire in our epistemic endeavors. The cost of endorsing Descartes' standard for knowledge is extremely high, and is

⁸³ As mentioned in Chapter 3, this phrase is borrowed from Keith DeRose, and articulates that requirement for defeating skepticism in which one provides a refutation that would knock the skeptic down.

particularly so when we consider that we have little motivation for endorsing this high standard for *every* kind of knowledge.

What I have aimed to show is that we do in fact possess evidence that grounds our ordinary beliefs; and Moore seems right to characterize that evidence as conclusive. In order to engage skepticism I have conceded to the skeptic that our sensory experience can only confer epistemic justification if the general assumption of reliability (AR) that underlies it is justified. So the debate on skepticism generally moves from our particular beliefs about the external world, to the more general question of the reliability of the senses. The skeptic wants us to try to show that we have grounds for thinking *sensory experience* in particular is reliable. However, assumptions of reliability for belief forming mechanisms seem to be grounded in the meta-principle AR rather than in any particular evidence that demonstrates the veracity of any particular method. Thus, I have shifted the *general* question of reliability that stems out of skepticism into a defense of AR as an appealing meta-principle in light of our pursuit of truth.

The sense in which Cartesian skepticism in the first person can be seen as a moral dilemma stems from the fact that the pursuit of truth and the reliance on sensory experience go hand in hand. Thus, the question about our reliance on sensory experience is really a question about whether or not the pursuit of truth justifies relying on a method that may or may not be reliable. Certainly an answer in the affirmative is an epistemic answer and does provide some epistemic justification in some sense of that term. However, we must admit that by relying on sensory experience we are also introducing not only the possibility of attaining truth, but also the possibility of falling into error. That is the dilemma, and that is the sense in which we can see fully that it is right to see

skepticism as a stalemate between the optimists on the one hand and the pessimist on the other.

While this final position may not be as satisfactory as we might like, it seems clear that it is the best possible response and that it is genuinely defensible in many ways. In light of the above arguments it seems, at least to my mind, that the skeptic's stance is no more defensible than the individual who refuses to run a race because in doing so she might lose rather than win. Such an attitude is no more admirable in the epistemic arena than it is in any other. It seems as though this first person question of skepticism boils down to an age old, common-sense truth that all things worth having involve both work and risk. In this light it seems surprising that philosophers ever thought that truth, as valuable as it is, is something we can and should attain without assuming any risk in the pursuit of it.

Bibliography

- Allen, Joshua. "Relevant Possibilities." *Philos Stud* (2008) 138: 55-71.
- Audi, Robert. *The Structure of Justification*. Cambridge: Cambridge University Press, 1993.
- Bonjour, Laurence and Ernest Sosa, *Epistemic Justification*. Oxford: Blackwell Publishing, 2003.
- Brueckner, Anthony. "Semantic Answers to Skepticism." *Skepticism: A Contemporary Reader*, edited by Keith DeRose and Ted A. Warfield, 43-61. Oxford: Oxford University Press, 1999.
- Brewer, Bill. "Externalism and A Priori Knowledge of Empirical Facts." *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, 384-414. Oxford: Oxford University Press, 2000.
- Cohen, Stewart. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 706-720. Malden: Blackwell Publishing, 2008.
- Davies, Martin. "Externalism and Armchair Knowledge." *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, 384-414. Oxford: Oxford University Press, 2000.
- DeRose, Keith. "Introduction: Responding to Skepticism." from *Skepticism: A Contemporary Reader*, edited by Keith DeRose and Ted A. Warfield, 1-26. Oxford: Oxford University Press, 1999.
- DeRose, Keith. "Solving the Skeptical Problem." *Skepticism: A Contemporary Reader*, edited by Keith DeRose and Ted A. Warfield, 183-219. Oxford: Oxford University Press, 1999.
- Descartes, Rene. "Meditation I." *Philosophical Skepticism*, edited by Charles Landesman and Robin Meeks, 45-50. Oxford: Blackwell Publishing, 2003.
- Dretske, Fred. "The Case Against Closure." *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 13-25. Oxford: Blackwell Publishing, 2005.
- Dretske, Fred. "Epistemic Operators." *Skepticism: A Contemporary Reader*, edited by Keith DeRose and Ted A. Warfield, 131-144. Oxford: Oxford University Press, 1999.
- Empiricus, Sextus. "Outlines of Pyrrhonism." *Philosophical Skepticism*, edited by Charles Landesman and Robin Meeks, 34-44. Oxford: Blackwell Publishing, 2003.
- Feldman, Richard and Earl Conee. "Evidentialism." *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 310-321. Malden: Blackwell Publishing, 2008.
- Foley, Richard. "Skepticism and Rationality." *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 322-332. Malden: Blackwell Publishing, 2008.
- Fumerton, Richard. *Metaepistemology and Skepticism*. Lanham: Rowman & Littlefield Publishers, Inc., 1995.

- Gettier, Edmund. "Is Justified True Belief Knowledge?." Readings in Epistemology, edited by Jack. S Crumley II, 131-133. Mountain View: Mayfield Publishing Company, 1999.
- Greco, John. Putting Skeptics in Their Place: the Nature of Skeptical Arguments and Their Role in Philosophical Inquiry. Cambridge: Cambridge University Press, 2000.
- Hume, David. "An Enquiry Concerning Human Understanding." 2nd edition, edited by Eric Steinberg, 1-114. Indianapolis: Hackett publishing, 1993.
- Klein, Peter. "How A Pyrrhonian Skeptic Might Respond." Epistemology: An Anthology, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 36-50. Malden: Blackwell Publishing, 2008.
- Knavig, Jonathan. "Contextualism, Contrastivism, Relevant Alternatives, and Closure." Philosophical Studies (2007) 134: 131-140.
- Kurtz, Paul. The New Skepticism. Buffalo: Prometheus Book, 1992.
- Lewis, David. "Elusive Knowledge." Skepticism: A Contemporary Reader, edited by Keith DeRose and Ted A. Warfield, 220-242. Oxford: Oxford University Press, 1999.
- Moore, G.E. from "Proof of an External World." Philosophical Skepticism, edited by Charles Landesman and Roblin Meeks, 125-13. Oxford: Blackwell Publishing, 2003.
- Nagel, Thomas. Selections from *The View from Nowhere*. Skepticism: A Contemporary Reader, edited by Keith DeRose and Ted A. Warfield, 272-291. Oxford: Oxford University Press, 1999.
- Pritchard, Duncan. Epistemic Luck. Oxford: Oxford University Press, 2005.
- Pritchard, Duncan. "Recent Work on Radical Skepticism." American Philosophical Quarterly (2002) 39.3: 215-253.
- Pritchard, Duncan. "Anti-Luck Epistemology." Synthese (2007) 158: 277-297.
- Putnam, Hilary. "Brains In a Vat." Skepticism: A Contemporary Reader, edited by Keith DeRose and Ted A. Warfield, 27-42. Oxford: Oxford University Press, 1999.
- Scaffer, Jonathan. "Knowledge, relevant alternatives, and missed clues." Analysis (2001) 61.3: 202-208.
- Sosa, Ernest. "The Raft and the Pyramid." Human Knowledge: Classical Contemporary Approaches, edited by Paul K. Moser and Arnold van Nat, 341-356. New York: Oxford University Press, 1995.
- Stanley, Jason. "Knowledge and Practical Interest." Selections in Epistemology: An Anthology, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 721-741. Malden: Blackwell Publishing, 2008.
- Stine, Gail. "Skepticism, Relevant Alternatives, and Deductive Closure." Skepticism: A Contemporary Reader, edited by Keith DeRose and Ted A. Warfield, 145-155. Oxford: Oxford University Press, 1999.
- Stroud, Barry. "Scepticism, 'Externalism', and the Goal of Epistemology." Skepticism: A Contemporary Reader, edited by Keith DeRose and Ted A. Warfield, 292-304. Oxford: Oxford University Press, 1999.
- Unger, Peter. "A Defense of Skepticism." Philosophical Skepticism, edited by Charles Landesman and Robin Meeks, 90-110. Oxford: Blackwell -Publishing, 2003.

- Warfield, Ted A. "A Priori Knowledge of the World: Knowing the World by Knowing Our Minds." *Skepticism: A Contemporary Reader*, edited by Keith DeRose and Ted A. Warfield, 76-92. Oxford: Oxford University Press, 1999.
- Williams, Michael. "Epistemological Realism." *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 51-72. Malden: Blackwell Publishing, 2008.
- Williamson, Timothy. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.

Vita
Jennifer A. Woodward
kuoppala@mscd.edu

Education

- Metropolitan State College of Denver** **1996-2000**
- Attained a BA in Philosophy with a minor in History in the Spring of 2000
- University of Kentucky** **2000- Present**
- Attained a Master's degree in Philosophy in December 2003

Employment

- University of Colorado at Boulder** *Part-time instructor* **Spring-Fall 2007**
- I am currently teaching an Introduction to Philosophy course.
- Metropolitan State College** *Part-time Instructor* **Jan. 2006-Present**
- I am currently teaching Introduction to Philosophy and Introduction to Ethics. I have taught critical thinking courses at Metro as well.
- University of Arizona** *Geography Research Assistantship* **Dec. 2005-
June 2005**
- I researched for requested information, data, photos, and maps on the internet and in the library.
- University of Kentucky** *Teaching Assistantship* **Aug. 2000- May 2004**
- As a graduate student I taught various stand alone sections of Introductory Philosophy courses, assisted in one section of Symbolic Logic, held office hours, and tutored students.

Service, Awards and Presentations

- *Dissertation Fellowship*, University of Kentucky, awarded for Fall semester 2004. I was nominated by the Faculty at University of Kentucky for this fellowship.
- *Teaching Assistanship*, University of Kentucky, awarded for each academic year from Fall 2000- Spring 2004.
- *Commentary Presented at the Southern Association for Philosophy and Psychology*, New Orleans, April 2004.
- *Served as co-chair for the 5th Annual Graduate Student Conference*, department of Philosophy at the University of Kentucky, 2003.
- *Served as Graduate student representative for the Chair Search Committee*, University of Kentucky, Spring 2002. I was nominated by the faculty for this appointment.
- *Served as Secretary on the Philosophy Graduate Student Association*, University of Kentucky, Fall 2001-Spring 2002.
- *Participated in the informal Reading Group in the Philosophy department*, University of Kentucky, each academic year Fall 2000- Spring 2004. This reading group provided an informal atmosphere for graduate students and faculty to discuss voted on philosophical works.
- *Participated in the Philosophy Department's Women's Forum*, University of Kentucky, Fall 2003- Spring 2004. This forum is dedicated to discussing and addressing issues of particular interest to women in Philosophy.