

Summer 2016

# Knowledge, infallibility, and skepticism

Gregory Douglas Stoutenburg  
*University of Iowa*

Copyright © 2016 Gregory Douglas Stoutenburg

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/6505>

---

## Recommended Citation

Stoutenburg, Gregory Douglas. "Knowledge, infallibility, and skepticism." PhD (Doctor of Philosophy) thesis, University of Iowa, 2016.  
<https://doi.org/10.17077/etd.pm1u43z0>

---

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

Part of the [Philosophy Commons](#)

KNOWLEDGE, INFALLIBILITY, AND SKEPTICISM

by

Gregory Douglas Stoutenburg

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

August 2016

Thesis Supervisor: Professor Richard A. Fumerton

Copyright by

GREGORY DOUGLAS STOUTENBURG

2016

All Rights Reserved

Graduate College  
The University of Iowa  
Iowa City, Iowa

CERTIFICATE OF APPROVAL

---

PH.D. THESIS

---

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

Gregory Douglas Stoutenburg

has been approved by the Examining Committee for  
the thesis requirement for the Doctor of Philosophy degree  
in Philosophy at the August 2016 graduation.

Thesis Committee:

---

Richard Fumerton, Thesis Supervisor

---

Evan Fales

---

Diane Jeske

---

David Cuning

---

Ali Hasan

In memory of my grandfather, Larry R. Demerest  
1932-2015

My habitual opinions keep coming back, and despite my wishes, they capture my belief, which is as it were bound over to them as a result of long occupation and the law of custom. I shall never get out of the habit of confidently assenting to these opinions, so long as I suppose them to be what in fact they are, namely highly probable opinions – opinions which, despite the fact that they are in a sense doubtful, as has just been shown, it is still much more reasonable to believe them than to deny. In view of this, I think it will be a good plan to turn my will in completely the opposite direction and deceive myself, by pretending for a time that these former opinions are utterly false and imaginary. I shall do this until the weight of preconceived opinion is counter-balanced and the distorting influence of habit no longer prevents my judgement from perceiving things correctly. In the meantime, I know that no danger or error will result from my plan, and that I cannot possibly go too far in my distrustful attitude. This is because the task now in hand does not involve action but merely the acquisition of knowledge.

René Descartes  
Meditations on First Philosophy

## ACKNOWLEDGEMENTS

Completing one's dissertation is a milestone, so I want to mention a number of people who helped me go that mile. Bernard Walker is largely responsible for my interest in philosophy. It is thanks to him that I first wanted to be a philosophy teacher and (I hope) stopped accepting easy justifications for beliefs I already held. James Swindal was extremely encouraging and generous with his time when I was probably pretty clearly a fish out of water struggling to find my philosophical home. Thanks to Walter Hopp, I did find my home: his epistemology class in the final semester of my studies for the Master's degree at Boston University convinced me to stay in philosophy when I very nearly left. When it was time to apply to PhD programs, he told me to apply to University of Washington and University of Iowa. The way I was thinking about epistemology must have been evident to him when it wasn't to me.

Each member of my dissertation committee—Evan, Diane, David, and Ali—has been extremely helpful and supportive, both in shaping the project and for my philosophical development. It was in their classes and through working with them while a teaching assistant that I learned how to think like a philosopher. Carrie Figdor helped me on a number of occasions by providing professional advice and extensive feedback on papers and ideas. It has truly been a privilege working with Richard Fumerton and learning how to be a philosopher from his example. Among the things I have learned from him is that doing serious epistemology involves doing serious systematic philosophy. As an advisor, Richard's dedication is incredible. I once had revisions due the following day for a paper he hadn't read (because, being a novice, I never showed it to him). He gave me comments by phone from the back of a taxi after reading my paper on a flight to visit

his family in New York. That showed a serious commitment to my professional well-being, and I am grateful both for that commitment and for his example.

I owe two substantial intellectual debts for ideas in the dissertation. The first is to the epistemologists on my committee: Richard, Evan, and Ali. Though I haven't met him, Panayot Butchvarov's epistemology has also been very influential to me. I'd like to think my dissertation fits comfortably within the Iowa epistemology tradition represented by these four. (How about that? Another person whose views largely agree with those of his most direct influences!) The second is to Peter Unger's book *Ignorance*, which is responsible for getting me to think about the importance of careful reflection on how we talk about knowledge and the significance of such reflection for substantive epistemological issues. I hope I am not committed to a view as radical as his.

My sharpest critics and closest allies have been my fellow graduate students in this philosophy department. I have received better feedback on papers during a Wednesday night graduate salon than from presenting the same work at multiple conferences. More often than not, those presentations turned into publications as a result of criticism from my peers. Bryan Appley deserves special recognition for never letting me get away with anything. Ryan Cobb has many times helped me turn rough ideas into slightly less rough ideas, and his cool demeanor often helped me find perspective. Landon Elkind is a gifted intellect, ready to talk through any idea. I've made regular use of that opportunity. Brian Collins also helped me talk through a number of ideas, most importantly by providing extensive feedback multiple times on what eventually turned into my first publication. Sam Taylor helped develop my sense of which epistemological issues are the crucial ones and gave me something to aspire to as a senior student in the

department. Brett Coppenger (along with Megan Coppenger) welcomed my family to the department and Brett showed by his example how to balance philosophical work and having a life. More than anyone else, Kris Phillips inspired me to give it my best shot over and over, whether it was turning in a conference paper, starting a philosophy summer camp, revising a rejected paper, or leading a discussion section. He has been my PhD coach.

I am extremely grateful to my family for their constant support. My grandmother, Patricia, made my graduate education possible. My brother, Drew, has been a source of level-headedness upon which I have often relied. My parents, Sheryl and Doug, are responsible for giving me the skills necessary to take on this past decade of academic challenges, professional challenges, marriage, and children—that whole collection of responsibilities that set the foundation for the rest of my life. My in-laws, George and Sandy, have been very supportive throughout my graduate school career. My wife, Heather, and my sons, Ulysses and Sebastian, are a source of immense joy, and it is due to Heather's selfless devotion that my graduate education is *finally* at its end. Not only that, but she's a sharp-minded philosopher herself. My work in this dissertation and elsewhere has been greatly improved by her criticism and insight.

This project is dedicated to my grandfather, Larry Demerest, who passed away just one semester before its completion.

## ABSTRACT

I argue that to know that a proposition is true one must have justification for being certain that the proposition is true. That is, one must have infallible epistemic justification for believing the proposition. It is widely accepted among epistemologists that we rarely, if ever, have such strong justification for our beliefs. It follows that there is precious little that we know. That conclusion is unacceptable to many philosophers. I argue that the positions that lead to the skeptical conclusion are well-motivated and that the skeptical conclusion is implicitly accepted by ordinary speakers.

My dissertation has three main components: a metaphilosophical position, an epistemological position, and an error theory. First, the metaphilosophical position. One very important part of philosophy is the analysis of our ordinary concepts. Analysis of our concepts begins with reflecting on what we are inclined to say about various actual and possible cases when considered under specific descriptions. This traditional method has recently come under attack, due in part to the rise of semantic externalism. I agree with externalists that if ‘meaning is reference’ then there is little reason to think reflecting on our concepts from the armchair will provide insight to the nature of the concepts we investigate. I defend a version of semantic internalism which grounds meaning in factors with which subjects are directly acquainted. That view supports the traditional methodology. Furthermore, as the goal of philosophical analysis is to accurately describe concepts of philosophical interest, the only kind of objection that could be decisive against a proposed analysis is that it does not correctly describe our concept. That opens the door to a skeptical analysis of knowledge.

Second, the epistemological position. I argue that the unacceptability of sentences of the form “S knows that p but it is possible for S that not-p” is best explained by the hypothesis that our concept of knowledge requires having justification for being certain that what one believes is true. I offer as a criterion of justified certainty the idea that when one knows a proposition is true, one is in a position to decisively answer questions about one’s knowledge of that proposition. I survey a number of competing theories of knowledge and show that they allow for the possibility of knowledge when one fails my criterion of decisive answerability. Those views fail my criterion because according to those views there is nothing the subject is aware of that guarantees for the subject that the allegedly known proposition is true. On that basis, I contend that knowledge is direct awareness of the factors that constitute the truth of the proposition one believes.

Third, the error theory. Of course, we rarely have direct awareness of the factors that constitute the truth of the propositions we believe. So, our knowledge attributions are generally false. Yet, they are overwhelmingly natural to make. I argue that competent speakers are often quick to recognize knowledge attributions as a kind of ‘loose talk’ akin to the way we loosely ascribe geometric properties to ordinary objects. We regularly call objects ‘square’ that are not even close to being square, and we are quick to recognize this when challenged. I argue that we do this because we are so accustomed to communicating using strict language to make false claims that we often do not attend to the literal falsity of what we are saying. While we accept very demanding standards for knowledge, the phenomenon of recognizing that our knowledge claims are regularly false is as pervasive as our use of loose speech.

## PUBLIC ABSTRACT

If you asked someone what time it is and that person simply told you, “It’s ten o’clock,” without so much as checking a watch, clock, or cell phone, you would probably think that person is simply guessing at what time it is. You surely would not think the person *knows* what time it is. And you probably think that the person does not know what time it is because the person does not have a good basis for reporting that it is ten o’clock. I think that is because when someone knows that something is true, that person is in a position to answer once-and-for-all questions about the grounds that give that person knowledge. The person who rattled off, “It’s ten o’clock,” does not know because that person has no basis for answering the question, “How do you know it’s ten o’clock?”

In these pages, I defend the view that when someone knows something is the case, that person is in a position to decisively answer questions about his or her knowledge. However, there is important resistance to this idea from two fronts. First, some philosophers have challenged the methodological assumptions that led us to this conclusion. Second, upon further reflection it becomes clear that if our conclusion is true, then we do not know as much as we often think we do. This dissertation defends its methodological assumptions against those challenges and explains that while our knowledge is rather limited, we are quick to recognize that fact.

## TABLE OF CONTENTS

PREFACE.....	xi
CHAPTER ONE: METAPHILOSOPHICAL ISSUES.....	1
1.1 The Central Metaphilosophical Issues Addressed in This Chapter.....	1
1.2 What Kind of Knowledge are We Analyzing?.....	5
1.3 Philosophical Analysis.....	7
1.4 Chapter Summary.....	63
CHAPTER TWO: INFALLIBILITY AND CERTAINTY.....	65
2.1 Knowledge Requires Certainty.....	66
2.2 In Search of Certainty.....	73
2.3 Chapter Summary and Conclusion.....	125
CHAPTER THREE: EVIDENCE, BASING, AND THE STRUCTURE OF KNOWLEDGE.....	128
3.1 Evidence.....	128
3.2 The Structure of Knowledge and Justification.....	135
3.3 The Extent of Knowledge and Justification.....	167
3.4 Chapter Conclusion: Skeptical Epistemology.....	172
CHAPTER FOUR: SKEPTICISM AND KNOWLEDGE ATTRIBUTIONS.....	176
4.1 The Problem of False Speech.....	177
4.2 Methodological Assumptions for the Analysis of “S knows that P”.....	198
4.3 Chapter Summary.....	231
AFTERWORD.....	233
REFERENCES.....	235

## PREFACE

This dissertation is an essay in epistemology in a traditional style. It shares its subject-matter with other works that have attempted to articulate the nature of core epistemic concepts like knowledge, justification, and rationality. Such works have a long history, going back at least to Plato's attempt in the *Theatetus* to defend the view that knowledge is true belief 'with an account'. Books and papers defending particular views on some core epistemic concept of course proliferate today, as they have for the past several decades. Not long ago, epistemologists would defend 'analyses' of the core epistemic concepts that served as their subject-matter. These analyses were supposed to be sets of individually necessary and jointly sufficient conditions. The rules of the traditional dialectic seemed to be getting the right intuitive verdict on hypothetical cases, avoiding regress, avoiding skepticism, and providing an informative account.

Then several things happened that disrupted the way traditional epistemological projects were pursued. Of course, plenty of epistemologists continued to pursue their traditional projects and offered responses to the movements I am about to mention, but there has been enough of a shift in many philosophers' thinking as a result of these movements that some discussion of the historical shifts is in order. First, Quine challenged the 'dogma' that there is a clear analytic/synthetic distinction in language. If there is no such distinction, then the project of offering *analytically* necessary and sufficient conditions seems to be simply misguided. This made philosophers worry about offering 'analyses' of epistemic concepts altogether. (History is a funny thing. I still can't come up with a halfway plausible argument that grants Quine's premises and concludes that philosophers can't offer analyses of whatever concept they want to analyze.)

Second, Gettier presented his cases that pointed to the need for a ‘fourth condition’ on knowledge beside justified true belief. A merry-go-round of increasingly complicated analyses of knowledge appeared to plug the leak, followed by criticisms of those analyses that revealed other leaks. After a while, this process seems to have led to massive fatigue, with some philosophers worrying again that knowledge may lack an analysis, or at least remaining deeply suspicious of further attempts to provide the needed fix.

Third, various forms of externalism arose to challenge the old way of searching for analyses. The first externalism argued that the meaning of many linguistic expressions depends in some crucial way upon the environment in which the expression is used. From there, it was argued that the meanings of many concepts in the mind also depend in some crucial way upon environmental factors. Around the same time, a distinctly epistemic externalism appeared with force. Epistemic externalists argued that more traditional analyses of epistemic concepts that impose conditions on knowers and justified believers to formulate arguments in defense of their beliefs, or come to know that their beliefs are formed on legitimate bases, or be able to appreciate the relevance of evidence to their beliefs, are all misguided. In place of the traditional approaches that placed heavy emphasis on what agents must be able to do in order to possess knowledge or epistemic justification, externalists proposed analyses that at least early on almost took the subject’s mind out of the picture altogether. A knower might just be someone whose belief was caused by its truth-maker, even if the knower is and would forevermore be oblivious to the belief’s relationship to what makes it true. The three externalisms—of language meaning, of mental content, and of epistemic concepts—collectively share the idea that the mind-

independent world is largely responsible for shaping the mind, against the older tradition that emphasized the subject's perspective in philosophy of mind and epistemology.

These three shifts in the way philosophers think about epistemic concepts are not clearly logically related. However, they occurred in such a sequence as to collectively cast suspicion on the traditional way of doing epistemology. To repeat, there have always been philosophers who resisted these developments, but altogether much philosophical thinking was still influenced by the developments just mentioned.

Separately, two other major changes took place. Particularly in the late nineties through the present, many epistemologists became less interested in the analysis of knowledge and more interested in semantics of knowledge attributions. The debate that began in the seventies over a relevant alternatives theory of knowledge received new life through epistemic contextualism, which, it is claimed, is a thesis about knowledge attributions, but not about knowledge itself. Of course, philosophers had for a very long time been motivating their analyses by means of reflecting on judgments that seem intuitive in response to thought experiments. What seemed different about this new way of doing epistemology is that it is so explicit about attempting to reach conclusions about the semantics (as if that's not just a fancy word for 'meaning'!) of some lexical item while often attempting to remain officially neutral about that which is expressed by the lexical item: some concept, property, or whatever.

Around the time of this development came the last one I'll mention. Many philosophers started getting worried that 'our intuitions' about this-or-that thought experiment may in fact be idiosyncratic. These philosophers proposed using the experimental methods of social science to attempt to get a better grasp of what 'our'

intuitions really are. (If some of these philosophers are right, ‘our’ intuitions are all over the place.) The movement is called experimental philosophy. Its proponents think experimental methods are essential—or if not essential, then at least very helpful—in defending epistemological claims these days, and they have been influential. It is no longer unusual to make a claim to the effect that such-and-such seems intuitive only to receive the objection that experimental data are needed to back the claim.

I suspect that the rise of emphasis on the semantics of knowledge attributions and on the sudden prominence of experimental philosophy are both related to the first three developments I mentioned, especially the rise of externalism(s). If correct answers to our questions about knowledge and justification are really to be found by looking out into the world, by getting out of our own minds, then it seems not implausible that surveying ordinary speakers about the correct way to use a term might actually provide some insight. And if the best we can do in discovering what knowledge *is* is to reach a kind of theoretical consensus as to when ordinary speakers do and do not predicate “knowledge,” then it makes a lot of sense to pursue the semantics of knowledge attributions over traditional epistemology; and if the inquiry into “knowledge” is this public matter, then of course we need to get into the minds of others to get to find the answers we seek. Still, the connections between these developments is far from clear: as I said, history is a funny thing.

The purpose of this (very rough) history of (a part of) epistemology is to explain why I have chosen to pursue this dissertation project in the way that I have. Given this history, I think any lengthy work in epistemology simply must defend its methodology against competitors, especially when one wants to pursue the project in a rather traditional

manner. I became convinced early in my doctoral studies that many of the controversies found in epistemology can be better understood by making clear the metaphilosophical and methodological assumptions that support the conflicting views. For instance, if one takes a pretty traditional view about mental content, then it should hardly be shocking when one comes to the conclusion that a large part of ordinary talk is false and often misleading. However, I am not so naïve as to think that substantive debates can be resolved by ‘moving up’ to argue about metaphilosophical assumptions instead. Many philosophers have thought that by finding a principled philosophical method they would find the answers they sought, only to be disappointed. Often enough, philosophical opponents share at least roughly the same views in other areas. Still, I think there are cases where theses in one area of philosophy fit much better with some further claims than others. Certain methodological claims are of this sort. I will explicitly argue in chapter four that the metaphilosophical views I favor support a skeptical view of knowledge attributions over other views in the debate.

In particular, a kind of internalism about meaning fits very well with a rejection of experimental methods in philosophy and with a rejection of the kind of ‘commonsense’ epistemology that characterizes the work of an Austin or a Chisholm. In addition, I think that same kind of internalism is the only fit for taking seriously the idea that we can arrive at answers to questions about the meanings of certain terms (or the analysis of certain concepts) by way of reflecting on what we are inclined to say about particular scenarios. Interestingly, however, if we hope to reach epistemological conclusions that are generalizable beyond the case of one’s own use of words or application of concepts, then what others find natural to say/predicate/think very well may be relevant to philosophical

inquiry. It is for that reason that while in the first chapter I argue that many of the trends mentioned earlier ought to be rejected, in the final chapter I defend a particular view about the semantics of knowledge attributions. I really do want to say that *we* accept very high standards for knowledge and that upon some reflection *we* accept that those standards are rarely met.

Each chapter in this dissertation has its own brief introduction. I won't preempt those introductions here. Instead, I'll indicate how the chapters connect. The first chapter focuses heavily on methodological issues, especially on defending the traditional method of reflection on hypothetical scenarios for the purpose of figuring out the contents of one's own concepts. I argue there that philosophical reflection relies heavily upon the use of language, so much so that reflection on language may be practically indispensable for philosophical analysis. In the second chapter, I use an argument centered on the intuition that sentences of the form "S knows that p but it is possible for S that not-p" are contradictory in order to motivate a demanding standard for knowledge, one that requires the elimination of all possibilities incompatible with the truth of the proposition one believes. In the third chapter, I spell out and defend some controversial consequences of my preferred analysis of knowledge. The view implies, among other things, that we know very little of what we regularly think and say we know. In the fourth chapter, I do two things. One is to argue for the superiority of a skeptical treatment of knowledge attributions. The other is to show that the superiority of a skeptical treatment of knowledge attributions provides the way for a highly plausible explanation of why we often say and think that we know a lot.

# **CHAPTER ONE:**

## **METAPHILOSOPHICAL ISSUES**

### *Chapter Preview*

In this chapter I lay out my favored positions on a number of preliminary issues relevant to the subject-matter of this dissertation. Here I identify the subject-matter of the dissertation and defend the method of thought experiment as a means to discovering the contents of concepts of philosophical interest. I think commitment to that method of philosophical analysis involves commitment on other substantive issues that I also discuss here, including the role of empirical work in philosophy, the semantic internalism/externalism controversy, and the relationship between ideas about kinds and necessity. My hope is that by getting a lot of foundational issues concerning philosophical method out of the way at the start, my later arguments for more straightforwardly epistemological theses will neither be burdened by methodological asides nor undermined by undefended metaphilosophical assumptions.

### **1.1 The Central Metaphilosophical Issues Addressed in This Chapter**

It has become clear to me while studying philosophy that commitments in one area often require commitments in others. In my case, some of the arguments I find most persuasive in epistemology ultimately require commitments in philosophy of language and philosophy of mind. Even where the connections are not so tight that one rationally *cannot* accept the conclusion of some argument without accepting the conclusion of another, there are many cases where the sort of argument that supports a particular conclusion at least requires defense in the form of premises from a different sub-field. As

an example, later in this dissertation I will challenge certain positions about the analysis of knowledge attributions—sentences of the form “S knows that p.” The most natural classification of such arguments is that they are epistemological in nature. But defense of my position there requires saying a lot about how speakers are sensitive to the meanings of the terms they employ, and that requires discussing issues in philosophy of language and philosophy of mind closely related to the semantic internalism/externalism controversy.<sup>1</sup>

The purpose of this chapter is to lay out the metaphilosophical positions I will assume in this project and to equip myself with philosophical tools to work on the problems I discuss. My hope is that I can offer a clear statement of what I take myself to be doing in order that when I consider various philosophical controversies I have something to lean on to support my position. In this way, my proposed solutions to problems will be principled rather than arbitrary. In addition, I will briefly defend my positions in philosophy of language and philosophy of mind that are needed to defend my epistemological views.

My project is to motivate and defend a certain analysis of the concept of knowledge. To many epistemologists, that very statement of the project likely sounds rather quaint and hopeless. One who shares this view might point to two important

---

<sup>1</sup> Some of the relevant issues in philosophy of language and philosophy of mind are themselves epistemological in character: a plausible take on a crucial metaphilosophical issue posed by the semantic internalism/externalism controversy is that some account has to be given of how speakers can come to know the meanings of the words they use. Clearly, issues about how one ‘comes to know’ anything are epistemological. The point here is that there are several complicated, related issues that ought to be addressed in a dissertation on the analysis of knowledge.

philosophical events. First, it has been several decades since Quine ‘showed’ that the analytic/synthetic distinction is unfounded. Second, and not long after the first, some philosophers argued very influentially that meanings—things likely to be among the best candidates for the things we attempt to analyze from the armchair—are often or usually out of our cognitive reach: they are not in the head, but out in the world, and the best way to study them is to engage in empirical work of the sort that philosophers have historically thought unnecessary.

So, before I start offering arguments for my preferred analysis of knowledge I think it is best to step back and lay some foundations. The statement of my project provokes questions. Does analysis consist in the examination of *concepts* rather than something else? If it does, shouldn’t this be a psychology dissertation? Surely psychologists are best equipped to tell us what concepts we have. Does analysis require commitment to *analytic sentences* of the sort whose existence Quine called into doubt? If knowledge is a real property, out there in the world, then how can you use thought experiments with a hope of delivering conclusions about knowledge—aren’t inferences from conceivability to possibility problematic?<sup>2</sup> Don’t many contemporary epistemologists claim “knowledge first,” arguing that knowledge is not analyzable?<sup>3</sup>

---

<sup>2</sup> The philosophers I have in mind here are the paradigm semantic externalists Saul Kripke and Hilary Putnam. Hilary Kornblith later argued that epistemological attention should be directed not to our *concept* of knowledge, but knowledge itself: a natural kind property like water. I engage with these arguments later in this chapter.

<sup>3</sup> Here I have in mind Timothy Williamson, whose defense of “knowledge first” epistemology has been very influential since he proposed it at the start of this century. I discuss Williamson in a few places, especially section 1.3.5.

I will address these sorts of worries in this chapter. Given the scope of the issues just raised, it will be impossible to give more than a sketch of my favored solutions to some of these problems. My goal here is not to decisively refute all influential positions incompatible with mine, but to attempt to offer a view that avoids a number of the problems that have been raised against traditional philosophical method in recent decades. The influence of some of the problems addressed here is so huge that a book-length project that ignores these concerns would be incomplete.

The chapter proceeds as follows. (1.2) After identifying propositional knowledge as the subject of the dissertation, I identify (1.3.1) concepts as the proper objects of philosophical investigation, and (1.3.2) defend thought experiments as the way to investigate concepts. Then (1.3.3) I discuss the semantic internalism/externalism controversy: if meanings are not ‘in the head,’ then it would seem that first-person reflection on hypothetical scenarios will be unable to deliver (on a rational basis) true philosophical conclusions. There, I offer two different versions of internalism about meaning and show how both are compatible with armchair analysis. Next, (1.3.4) I consider a few further issues concerning philosophical analysis: specifically the analytic/synthetic distinction, whether there are any analytic sentences, whether and how any justification or knowledge is *a priori*, and some reflections on the nature of necessity as it relates to our topics of concern. Finally, (1.3.5) I conclude by addressing Timothy Williamson’s influential view that knowledge, in particular, is not analyzable. I argue that even if knowledge does not decompose into constituent concepts, it is still analyzable in an important sense.

## 1.2 What Kind of Knowledge are We Analyzing?

These sentences all use “know” appropriately.

- (1) “I know that I have hands.”
- (2) “I know you’ll do great! Break a leg!”
- (3) “I know Maggie. She’s a great physician.”
- (4) “I know how to get to the trailhead. Take a left at Main, walk a mile, then...”
- (5) “I know how to ski. Watch me!”

(3) states that the speaker bears a relation dubbed ‘knowing’ to a *person* and is able to state her qualifications as a physician. Similarly, (1) and (2) state that the speaker bears another relation picked out by “know,” but the relation in those cases is to a *proposition*: in (1) the proposition is *that I have hands* while in (2) the proposition is *that you will do great*.

All of these contrast with (4) which states that the speaker possesses a certain sort of ability: knowledge-*how* to perform the action of getting to the trail. Whether knowledge-*how* claims should be identified with knowledge-*that* claims is controversial (Ryle 1949, Stanley and Williamson 2001). It’s not terribly implausible that while (4) *superficially* makes a claim about knowing how to do something, it should properly be read as “I know *that* the trailhead is down Main, past the fire house,” in which case (4) would be a disguised use of an expression like (1). (5) is tougher. If I ask a friend, “Do you know how to ski?” and my friend says, “Yes,” I may not expect this friend to be able to explain to me how to ski. Rather, I’m wondering whether my friend would successfully navigate the terrain if given skis and pushed off of a lift at the top of the

mountain. Knowing *how* to ski, if it does reduce to a kind of knowing *that* the way to ski is X, does so in a complicated way.

A plain interpretation of these five sentences suggests that we have identified three types of knowledge, all of which we regularly attribute to ourselves and others by the use of “know” and “knowledge.” Arguments can be made, of course, that there are more or fewer types of knowledge than I have acknowledged. One might argue that perhaps (2) here, while mimicking the form of (1), is a disguised exhortation rather than a truth-evaluable claim—perhaps equivalent in meaning to “Go get ‘em!” (Chrisman 2007). I already mentioned the possibility that knowing *how* to X is identical with knowing *that* the way to X is by taking steps ABC. Whether knowledge-how turns out to be *sui generis* or reducible to knowledge-that, the examples we have considered suggest a distinction between knowledge *that* such-and-such is the case and other sorts of knowledge. Knowledge-that—the sort of knowledge that relates a knower to a proposition—will be my focus. This sort of knowledge is called *propositional* knowledge. Thinkers wanting to know how things are want to know what is true. The only kind of knowledge we introduced a moment ago that has a connection to truth is propositional knowledge. If the other sorts of knowledge do not reduce to propositional knowledge, one could know how to ski without that consisting in knowing truths about skiing, and one could know Maggie in the sense of having met her without thereby knowing truths about her. When we know that such-and-such is the case, such-and-such *is* the case, and we are related to it in the way that satisfies intellectual curiosity. There may be interesting connections between propositional knowledge and these other sorts. I

will not explore those connections. Rather, I will focus on the nature of propositional knowledge itself.

The project is to *analyze* propositional knowledge. But what is philosophical analysis? I now turn to that question.

### 1.3 Philosophical Analysis

Philosophers have for a long time been interested in answering questions about the nature of concepts of philosophical interest like goodness, knowledge, and right action. When they ask these questions they are after what is *essential* to goodness, knowledge, right action, and so on. They want to know what *has to be* the case for an action to be right, not what is coincidentally always the case or what is usually the case. The philosopher who offers an analysis means to be getting to the core of the object of analysis.

A philosophical analysis, traditionally, is an identification of one concept with other concepts more fundamental than it. An analysis of this sort takes the form of an identity statement, e.g. “knowledge *is* justified true belief.” To *search for* an analysis is to engage in a project of attempting to understand systematic connections between one concept of philosophical interest and others. A philosopher might search for an analysis of knowledge and turn up absolutely nothing in the way of statements that successfully identify the concept of knowledge with more fundamental concepts. What do we say here? I think we should say the philosopher’s search for an analysis of that concept did not yield an analysis of the concept. Still, much philosophical work consists in the activity of *searching* for satisfactory analyses of concepts of philosophical interest. It is for that reason that we can make perfect sense of the claim that Timothy Williamson

works on the analysis of knowledge but denies that knowledge has an analysis (2000) (see section 1.3.5): Williamson's search for an analysis did not result in an analysis.

It would be a mistake to claim that philosophy is nothing but analysis—to analyze *philosophy* as *conceptual analysis*—so that someone who claims to be doing philosophy but is interested in empirical methods, or spends a lot of time doing phenomenology, or whatever, is not really doing philosophy, despite what he says. And it would be a mistake not just because it's a bad way to make friends in one's department. While I think providing analyses is a paradigmatically philosophical activity, the paradigm does not define the borders of philosophy. Philosophers often rely on their own introspection of their intuitions about thought experiments, and to introspect is not to offer an analysis. Indeed, I think the idea of providing a necessary condition on doing philosophy is probably misguided, even while I think I can identify a sufficient condition, namely engaging in analysis. Understood in a very broad sense of 'engaging in analysis', though, philosophy probably is identical with engaging in analysis, for these activities that philosophers have occasionally disparaged—phenomenology, empirical work, etc.—can be relevant to the truth-value of a proposed analysis. Surely much of the reason I think the analysis of C is AB is because I was driven to that conclusion by thinking carefully about my own intuitions about thought experiments, and it would be very strange to say I had been doing philosophy when I entertained the thought "Is C identical to AB?" then briefly stopped doing philosophy to run a thought experiment, then resumed doing philosophy when I concluded "C is AB."

That concludes the ecumenical portion of this chapter. From here, I need to make a lot of commitments to put flesh on the bones of this conception of analysis and doing so will require taking sides on a number of issues.

### 1.3.1 *Objects of Analysis*

In the previous section I wrote of analyzing *concepts*. That is the view I want to defend. But while philosophers have often spoken of “conceptual analysis” and “analyzing the concept” of this-or-that, it has not always been clear what they take the target of their analysis to be. To describe philosophical analysis as conceptual analysis really is not as innocent as it appears to be unless the occurrence of ‘conceptual’ in “conceptual analysis” is inert. Some argument has to be made for why we should think of concepts as the objects of analysis. It is difficult to offer arguments with uncontroversial premises for theses of such a fundamental variety. I think, though, that an appeal to what it takes to be a competent speaker of a language can take us quite a distance toward understanding philosophy as the analysis of concepts.

When a philosopher performs an analysis there is something toward which the philosopher’s attention is directed, whether directly or indirectly. Perhaps the analysandum is something of which the analyst is directly aware, as Moore seemed to think (Moore 1903 (2000)). Or maybe it is partly hidden from the mind so that more indirect investigatory methods are needed. Whichever of these is the case, I think it is obvious that there is an existence requirement on the analysandum. If the philosopher is

analyzing X, then X exists.<sup>4</sup> It would simply be nonsense for a philosopher to claim to have been studying (say) the property *being a right action* only to conclude at the end of the investigation that there is no such property.<sup>5</sup>

That alone might seem sufficient to dispense with a whole slew of candidates for the analysanda: if there is no X, the thought goes, one can't be analyzing the property *being X*, or states of affairs that include X, or facts that include X. We can distinguish, though, between the *target* of analysis—what I have been calling the analysandum—and the *means* of analysis. The moral philosopher might claim to be interested in the property *being a right action*, but admit that one must provide substantial argument for the claim that there is any such property rather than taking it for granted as a condition of beginning a philosophical investigation. While the goal of analysis is to dissect the property of being a right action, something goes proxy for the property of being a right action: something whose existence is less controversial than the property of being a right action stands in for the duration of the investigation. This proxy is what I call the *means* of analysis. The moral philosopher equipped with this distinction is prepared to head off moral skeptics who claim the object of analysis is not the *property* whose existence is controversial: instead, perhaps, it is some other entity sufficiently similar to or constitutive of the property of being a right action.

---

<sup>4</sup> An existence requirement like this is present in Fumerton (1983, p. 480).

<sup>5</sup> Here I am reading “studying” as factive. It would not be nonsense to claim to be investigating whether or not there is a property *being a right action*. But to even consider that the question, “Is there a property *being a right action*?” might be answered negatively, we must grant that the object our attention is directed toward when we are engaged in analysis is not the property *being a right action*—for then it would exist—but something that goes proxy for it.

I think the distinction between a target of philosophical analysis and a means of analyzing the target is a sound one. It respects the existence requirement on objects of analysis and avoids begging any questions in favor of the existence of a particular target of analysis. But I think that once the distinction is made there is no longer much to be said in favor of taking properties, states of affairs, or facts as the targets of analysis. There has to be something present to the mind, directly or indirectly, upon which the philosopher's attention is directed in performing an analysis. And I can't think of a single plausible candidate for a property, fact, or state of affairs that can serve as a means of analysis without being identical to the analysandum. What property could be held before the mind and analyzed that could be used to understand the nature of right action without guaranteeing that there is a property of being a right action? And even if someone comes up with such a property, it is likely to be itself a controversial item in the moral philosopher's ontology—perhaps the property *being good*. Clearly appeal to that property does no better at respecting the possibility that the properties philosophers are interested in may not exist.

That's not a knockdown argument. But I think it effectively casts doubt on the idea that the entities just mentioned are the targets of philosophical analysis. To go the rest of the way, I'll make a phenomenological appeal. What seems to be going on when I perform a thought experiment is that I entertain some scenario under a certain description and ask myself whether that scenario is a case of F. I imagine an ailing grandmother on her deathbed handing to her grandson a check for \$50,000 made out to Cash with a single request: to give the money to her favorite charity, Water the Children. (They provide

resources to obtain freshwater in communities where it is hard to come by. A contentious committee meeting resulted in the less-than-apt name.) “Ok, grandma, I promise to give this money to Water the Children in your name.” The grandmother passes away. The grandson goes home and looks up Water the Children and discovers that 75% of funds donated go to administrative costs and refreshments for contentious committee meetings. He then discovers a charity with exactly the same purpose but where only 15% of donations go to overhead costs. He feels torn, but donates the money to this other charity. Now I ask myself, did he do the right thing by breaking his promise and giving the money to this alternative charity?

The way this scenario is described, there are *hypothetical* states of affairs, properties, facts, and whatnot that I entertain when I consider whether the scenario is an instance of right action. But the purpose they serve is to bring out my standard for right action. The answer to the question, “Did he do the right thing?” depends upon the way I’ve thought of the scenario and whether that description fits the standard of right action I accept. Whatever the targets of philosophical analysis are, the *means* of analysis have to consist in something elicited by thought experiments. I think *concepts* are the most natural candidates for means of analysis in the technical sense of ‘means’ I introduced above. One must have a concept of right action in order for that concept to be elicited by a thought experiment. And whether or not there is, in addition to my concept of right action, the property *being a right action*, it seems to be inert in my judgment that a particular scenario fits my concept. So whatever entities are admitted into one’s ontology, concepts are necessary for coming to conclusions about thought experiments,

and once the description of the scenario is understood, little else is needed besides the concept to come to a judgment about the thought experiment.

Analysis of concepts is indispensable for analyzing whatever it is one might be interested in analyzing. Perhaps a philosopher does want to know whether some particular property, fact, or state of affairs exists. Analysis of any of these things proceeds through analysis of concepts. That is what I mean by saying concepts are the means of analysis. It may also be that concepts are the targets of analysis: that philosophical analysis proceeds through analysis of concepts *because* concepts are the ultimate analyzanda of philosophy. For example, philosophers have long debated about whether or not God exists. To begin to answer the question, “Does God exist?” one must first have an understanding of what is meant by “God” in that question. Providing an analysis of the concept *God* is an indispensable first step toward answering the question, “Does God exist?”

One might argue that a necessary condition of possessing a concept is that some other entity E exist, where E is one of the candidates of analysis I have been trying to steer away from. At one point, Russell held that thoughts got their content by including universals as constituents (1912, p.52). So maybe one cannot have thoughts about such-and-such being a right action without the universal *being a right action* being involved in the thought. Perhaps similar arguments can be constructed for the same purpose, namely to argue that even though concepts are the means of analysis, concepts still require the existence of one or more of these other kinds of entities.

It would take us far off course to evaluate every claim about the analysis of concepts. In what follows I'll assume as neutral a characterization of concepts as possible. I want to note that I haven't refuted the claim that entities like properties, facts, and states of affairs might in some way be necessary for philosophical analysis that proceeds by way of considering whether or not a particular concept applies to a hypothetical scenario. But I think that I've managed to focus attention on what is most important in the evaluation of such scenarios, and that is the concept one employs in judging whether the scenario falls under a particular description. Debates about the nature of concepts and what entities must exist for one to possess a given concept are secondary.

### 1.3.2 *Concepts and Philosophical Analysis*

Philosophical analysis consists in analyzing concepts of philosophical interest. Concepts are constituents of thoughts. Different individuals can share concept (-types) and often do.<sup>6</sup> We assume when we do philosophy that we are not idiosyncratic, that our analyses of the concept of F apply just as well to the concept of F possessed by others. So, when we do philosophical analysis, we attempt to make explicit the way we think about things.

In the conclusion of the last section I said we can leave aside questions about the nature of concepts. While I think I can remain noncommittal about what *gives* concepts

---

<sup>6</sup> For my purposes, it is not necessary to defend a position about the metaphysics of concepts. If realism is true, then subjects grasp one and the same concept; if nominalism is true, then subjects entertain token-distinct but type-identical concepts. Either view can accommodate the idea that different subjects can entertain the same *contents*. I will have more to say about the nature of concept-sharing below. I am prepared to grant that the concepts we 'share' are not precisely identical, but are similar enough to allow us to communicate effectively and be said to think 'the same thing.'

content and what concepts are, I do have to say something about how we figure out the contents of our concepts. There's usually no direct way of figuring those contents out. I don't usually just look in on my concepts and "see" their content. The way I figure out what my own concept of C is is to reflect on what standards I hold for when something falls under C. That method is indirect. In fact, I think it's just the method of thought experiment: entertaining a hypothetical scenario and thinking about how I would classify the scenario in certain relevant ways. I can think about what my concept of knowledge includes by entertaining a scenario in which someone has a true belief formed as a matter of luck and asking myself whether I would consider that knowledge. (I wouldn't.)

One occasionally finds a quotation from a philosopher that suggests the view that philosophical analysis chiefly involves the analysis of the meaning of words, while holding similar views about the role of thought experiments in figuring out those meanings.<sup>7</sup> To those philosophers, we figure out the content of our concepts by figuring out how we would *describe* certain scenarios. Description requires language. Why not adopt this fully language-centric view? Because concepts are more fundamental than the words that express those concepts. Two people who speak different languages may well possess the same concept (-type) while they express it with different words. But once one recognizes the connection between a concept and a term that expresses it, one can describe the analysis of a certain *concept* as the analysis of the meaning of a certain *word* without much remainder.

---

<sup>7</sup> For a discussion of this sort of view, but not an endorsement, see Richard Hare's excellent paper (Hare 1960).

A certain view of word meaning may seem to be in a stronger position to resist this argument. On this view of word meaning, what I just said about the relationship between concepts and words is false. What I just said is that words express concepts and as a result thinking about how one would use a word often just is thinking about concepts. The view of word meaning that challenges my claim is a functionalist view of words. I find it in (Sellars 1963). Sellars says the ‘abstract entities’ of philosophical interest—Platonic forms, relations, that kind of thing—are linguistic expressions of a certain sort. His understanding of ‘linguistic expression’ is technical. The ‘abstract entity’ redness, on this view, is the linguistic expression named<sup>8</sup> by both ·red· and ·rouge·:

Redness...is the word ·red· construed as a linguistic kind or sort which is capable of realization or embodiment in different linguistic materials, e.g. ·red·, ·rot·, ·rouge·, to become the English word “red,” the German word “rot,” and the French word “rouge” (627-8).

So, on this view, ‘linguistic expressions’ are abstract entities. Here is how this view poses a challenge. On this functionalist view of words, a word includes a ‘linguistic expression’ in virtue of the inscribed or uttered token being a realizer of a ‘linguistic expression’. (I use scare quotes around ‘linguistic expression’ to highlight the abstract nature of these things.)

I think, actually, that this view of word meaning does not threaten what I said about the role of language in conceptual analysis, though it may complicate it. I already granted that we can analyze our concepts by thinking about how we would describe

---

<sup>8</sup> On Sellars’ view, naming is the relation between a dot-quoted expression and the linguistic type it picks out: “I shall use dot quotes to form the names of the expressions...which is realized in English by the sign design illustrated between them” (Sellars 1963, p. 627 fn. 3).

hypothetical scenarios using language. The objection says when we do philosophical analysis we never really leave language behind because the meanings analyzed are not to be identified with concepts, but with ‘linguistic expressions’. In defending the view that concepts are the things we analyze, I have assumed that what we discover when we consider our predicative dispositions are semantic values. Now, clearly, a defender of a Sellarsian view of words does not want to claim that we discover our uses of words in the sense of discovering our dispositions to deploy tokens of certain sound- or inscription-types. Rather, the Sellarsian view has us considering how we would deploy ‘linguistic expressions’ in the sense of those abstract linguistic types named by inscription-tokens surrounded by dot-quotes. These are the semantic values analyzed by philosophers.

I see little reason to refrain from claiming that our considered dispositions to categorize things using those semantic values named by dot-quoted expressions are either identical to concepts or still rely on concepts (but not because they are identical). It seems to me that I am capable of entertaining a hypothetical scenario and thinking, “That’s a case of knowledge” without thinking, “I would call that ‘knowledge’” or “That’s a case of ·knowledge·”. But perhaps I am not so capable as I think, and when I consider hypothetical scenarios I really am always predicating dot-quoted ‘expressions’, and thus analysis of language in this technical sense of ‘language’ is indispensable for philosophical analysis. I am comfortable noting the complication and moving forward: concepts may be identified with dispositions to predicate ‘linguistic expressions’.

Another important role played by language in philosophical analysis is in forming and communicating the content of concepts. I’m pretty sure that what I think of as

knowledge is largely if not entirely the result of the competence I've developed with "knowledge" as a speaker of English. Competent speakers say they know the President's name, where they'll be this summer, what time it is; they deny that anyone with a lottery ticket knows the outcome of the lottery, what other people are thinking (often), and who the next President will be. Competent speakers respond to challenges to their knowledge claims by retracting their earlier claims, by defending them, or by modifying their earlier claims in some way. Much more could be said about typical patterns of use, of course. Exposure to these patterns of use forms a new speaker's concept of knowledge.

All of the data relevant to a philosophical analysis of a concept expressed in language are available to any competent speaker.<sup>9</sup> As a result, all the information necessary for arriving at a philosophical analysis of a concept is available to anyone who is competent with both the language in which the term expressing the concept is employed and with the term itself. The philosopher interested in the analysis of knowledge attempts to understand how the concept of knowledge is employed in these patterns of use. Patterns of ordinary use are data for the analysis. Additional data are gathered by the use of thought experiments. As a result, patterns of ordinary use are not trusted not to mislead: it may turn out that certain thought experiments strongly suggest that the natural inclination to say someone comes to *know* what time it is by looking at a

---

<sup>9</sup> While the data are available, it's uncommon for a philosophically-untrained speaker to easily pick up on the relevance of the data to the analysis of the concepts they so routinely deploy. So, my use of "available" in the footnoted sentence should not be confused with the claim that competent speakers easily make use of that data for their own analyses.

clock is misled and thus “knowledge” is not the correct way to describe whatever relation that person stands in to the proposition describing the time.

One might at this point wonder what role should be played by psychology in philosophical analysis. Concepts, after all, are plausibly mental entities. Psychology is the study of the mind. So, perhaps there is no better place to look for philosophical insight than psychology. Goldman and Pust agree with me that philosophical analysis is the analysis of concepts, and they infer largely on that basis that philosophical analysis can benefit substantially from psychology (1998).

I think they are wrong, and for instructive reasons. One argument they offer is this. Some versions of epistemic contextualism say that whether or not some knowledge attribution is true depends upon what ‘relevant alternatives’ are entertained by the knowledge attributor. So, whether “S knows that p” is true depends upon what not-p possibilities are entertained by the speaker of “S knows that p.” If the knowledge attributor is thinking about brains in vats, then almost never will they attribute knowledge to someone of some proposition. But if the attributor is only thinking about more ordinary possibilities of error, they are often more likely to attribute knowledge, even knowledge of the very same proposition p.

Goldman and Pust continue by pointing to studies in psychology that suggest people answer questions differently about how well off they and others are depending on whether they’d been primed to think about someone who is very well-off versus someone who is less well-off. The gist is that what we’re inclined to say about different

propositions depends in part upon what possibilities we are entertaining at the time we consider the proposition. Here is the alleged philosophical upshot.

If an epistemologist is trying to choose from among competing theories, some of which make heavy appeal to attributor context and some of which do not, these psychological findings are highly relevant. They give greater plausibility to theories assigning a large role to attributor context than those theories would have in the absence of such findings. It is much more plausible to hypothesize that intuitions about knowledge are heavily context-driven if other sorts of judgments and evaluations are shown to be so driven (1998, p. 193).

Epistemologists have long been aware from their own experiences as ordinary speakers that people *use* “S knows that p” in different ways in different contexts: that knowledge attributions are, in a sense, rarer in epistemology class than they are on the street. No one needed a study about psychological contrast-effects to be convinced of that. We already had the relevant data.

Goldman and Pust’s appeal to the study concealed two significant claims. The first significant claim is that the psychological fact that knowledge attributions vary with contextual shifts supports the philosophical claim that the *truth-values* of knowledge attributions vary with contextual shifts. But the data do not *and cannot* support that. It is up to the philosopher to argue for the best analysis of the meaning of “S knows that p” with the psychological data in mind. There are many options consistent with the use-data. Maybe the best analysis is a contextualist view of the sort Goldman and Pust mention. Or maybe fewer knowledge attributions are true, but due to conversational norms it is appropriate for speakers to assert “S knows that p” anyway (Black 2005, Rysiew 2001). Or maybe no positive knowledge attributions are true yet conversational

norms still make it appropriate to make knowledge attributions.<sup>10</sup> Or maybe positive knowledge attributions are always false and it is, contrary to appearances, always in some sense *inappropriate* to attribute knowledge. The point here is that the psychological data are compatible with all of these interpretations of the meaning of “S knows that p.” To draw a conclusion about the meaning of that expression on the basis of use-data is a distinctly philosophical theory-constructing exercise, one that must be conducted from the armchair. The psychological data severely underdetermines the choice of analyses and consequently doesn’t support any of them to the exclusion of the others.

The second significant claim is that the version of relevant alternatives contextualism being discussed holds that the truth-value of “S knows that p” depends upon the not-p alternatives being *considered* by the knowledge-attributor. That’s an analysis of knowledge, the evaluation of which is not helped by appeal to psychological data. The philosopher’s question is not whether speakers *do* vary in their attributions of knowledge depending on what they’re thinking about. The philosopher’s question is whether our concept of knowledge counts someone as knowing that p in one circumstance and not knowing that p in another depending upon what possibilities are entertained by the attributor. We evaluate that proposed analysis in the usual way: by thought experiment. Consider whether we would judge true “S knows that p” in the following hypothetical circumstances. “P” is “there is coffee in my cup”; “S” is an ordinary person who drinks a lot of coffee and so is familiar with its taste and smell; and

---

<sup>10</sup> Davis (2006) suggests something like this, but it is unclear whether his view is as skeptical as the one I mentioned above.

there is a poisoner in the area who has been going around and replacing the black coffee in people's cups with the tea from a poisonous plant that looks and smells very much like black coffee, enough so that someone who isn't actively examining their coffee cup would be very unlikely to detect a difference in taste and smell. S is next on the poisoner's list.

Our knowledge-attributor—call her Kay—is completely unaware of the existence of the poisoner and so would not evaluate knowledge claims by contrasting p with the poisoner scenario. If asked, Kay would likely *say*, of any S who is now drinking something that appears to be coffee, “S knows that he is drinking coffee.” Let's assume the poisoner hasn't gotten to S yet, so the beverage in S's mug is in fact coffee. I'm inclined to say that token of “S knows that he is drinking coffee” is false: although he is drinking coffee, he doesn't know it. For all he knows he could be drinking poisoned tea. The fact that Kay doesn't *consider* that possibility does not turn S's true belief into knowledge. The range of relevant alternatives include possibilities *not* considered by attributors.

We figured out from the armchair that “S knows that p” can be false while p is true and no error-possibilities entertained by the speaker obtain. So, the analysis of knowledge Goldman and Pust drew from the psychological data is false. Furthermore, the falsity of “S knows that p” in the scenario is indeterminate between at least two competing theories of knowledge. On one, S doesn't know that p because there is *in fact* a poisoner who very nearly replaced S's coffee with poisoned tea. On the other, S doesn't know that p because S was unable to eliminate the poisoned tea possibility, so the

*actual* presence of the poisoner did nothing to further undermine S's epistemic status: S wouldn't have known there was coffee in the cup even if there had been coffee in the cup and no poisoner nearby. Philosophers working to support one of those theories over the other will have to appeal to yet further thought experiments, and so to data that remains out of reach of psychology.

While philosophical analysis is the analysis of concepts, the sort of data that are relevant is available to anyone who is a competent speaker, in the sense of 'available' I described above: the data are there for anyone to take notice, even though it is unlikely that philosophically naïve speakers will easily appreciate the relevance of the data to the analysis of concepts of philosophical interest. On the plausible assumption that philosophers are competent speakers, we already have available all the data psychology could provide for our analyses. The only real worry psychological studies of this sort might present is the possibility that philosopher's intuitions are somehow corrupt: maybe considering too many thought experiments and being enamored by one's own analyses slants one's view about what is and is not natural to say and think. While I think that concern should prompt philosophers to try as hard as possible to ignore their favored views when considering thought experiments and other philosophers' work, it does not suggest that psychology can lead us to see things correctly. How, after all, might psychology show that philosophers' intuitions are out of step with those of ordinary people? A popular approach of psychologists is to ask undergraduates how they would answer certain questions. It is obviously true that philosophically naïve undergraduates' intuitions are less likely than ours to be 'corrupted' by philosophy. But with that benefit

comes a drawback: they are far less equipped than philosophers to pay attention to or even understand subtle cues that are crucial to the evaluation of a scenario.<sup>11</sup> It would be nice if we could be *sure* that we are unbiased. We just can't be, with or without psychology. And since we have reason to believe we are competent speakers, able to distinguish what is and is not relevant to evaluating a thought experiment, the best we can do is rely on our own judgments (cf. Jackson 1998, p. 37-38).

Given my insistence on speaker competence providing the data philosophers need to do analysis, one might think that instead of looking to psychology for insight into our concepts, we should look to lexicographers for insight into our language. Lexicographers engage in extensive study for the purpose of writing dictionaries. Their work includes defining words, but extends far beyond it: recording the use of prefixes and suffixes, word equivalents across languages, collecting and organizing pronunciations, and plenty more. Gilbert Ryle reports that upon stating that philosophy “examines the meaning of expressions,” H.J. Paton asked him, “Ah, Ryle, how *exactly* do you distinguish between philosophy and lexicography?” (Kornblith 2002, p. 162, Ryle 1970). If we want to know what our terms mean, where better to look than lexicography?

There is something to this question. But before saying what it is, it takes no more than a glance at a dictionary entry for “knowledge” to see that philosophers are up to

---

<sup>11</sup> To take just one example, one prominent experimental philosophy paper relies explicitly on asking undergraduate University of Madison-Wisconsin students questions about whether a person in a story *knows* a proposition and also if that person *believes* the proposition (Myers-Schulz and Schwitzgebel 2013). They were presented cases like Colin Radford's case of the unconfident examinee: a student is asked a question, hems and haws, then gives the correct answer—the moral being that the student knew the answer but didn't believe it (1966).

something *very* different from lexicography. The Oxford English Dictionary lists plenty of uses of the term. Here are a few: “recognizing something already known about,” “the fact or condition of knowing something,” “acquaintance, familiarity gained by experience,” “sexual intimacy,” “the faculty of understanding or knowing,” “the apprehension of fact or truth with the mind; clear and certain perception of fact or truth; the state or condition of knowing fact or truth.” While some of these elucidations begin to offer some insight into what knowledge is, most of them are circular, including “knowledge” or a cognate in the entry on “knowledge.” Surely philosophers would not let that pass for an analysis.

But the circularity of dictionary elucidations and examples is not the primary thing that distinguishes them from philosophical analyses. The lexicographer studies *actual* use and use alone and looks for similar features across various patterns of use. That patterns vary is partly responsible for the diversity of uses of “knowledge.” Philosophers seek a unified *definition* even if they are interested, as Ryle was, in the study of expressions. We want to know what underlies all of these ‘senses’ of “knowledge” and we won’t be content merely to offer several distinct patterns of use or disjoint analyses simply because the word is used in a variety of ways. We want to see what it is that *makes* similar uses similar. Merely identifying similarities is not enough.

To find what unifies different patterns of use requires avoiding heavy emphasis on *actual* use. I don’t mean to suggest actual use is irrelevant. Rather, as I have been stressing, we seek to understand the concept we express with a term rather than compiling our actual uses of the term and inferring the concept from that. Our study

requires including how we *would* classify hypothetical scenarios under specific descriptions. Thus the terms we employ in ordinary language might not *mean*, in the philosopher's sense, what they *mean* in the lexicographer's sense. For that reason, it is a little misleading to say philosophical analysis is the study of the meaning of expressions. The philosopher will understand this study of meaning to be of a different nature than the kind of study of meaning that occupies the lexicographer.

### 1.3.3 *Semantic Internalism, Semantic Externalism, and Meaning Analysis*

All this talk about armchair meaning analysis invites a discussion of the semantic internalism/externalism controversy, for the currently popular semantic externalist view at least appears to be in tension with the idea that one can accurately describe the meanings of terms in one's own language from the armchair.<sup>12</sup> Indeed, the position on philosophical analysis I have so far defended presupposes a kind of semantic internalism. The literature on this debate is enormous. I'll ignore many complications and restrict myself to addressing some key issues and sketching the position I find most plausible.

In his landmark 1973 paper "Meaning and Reference," subsequently expanded in his 1975 paper "The Meaning of 'Meaning,'" Hilary Putnam famously argued that meanings "ain't in the head." He asked us to consider a certain thought experiment. Putnam claimed that in 1750, before it was discovered that the chemical structure of

---

<sup>12</sup> For what it's worth, in a recent survey of philosophers—specifically, professional philosophers, holders of PhD degrees in philosophy, and graduate students—51.1% answered "accept or lean toward externalism" when asked their view of mental content. 20.0% answered "accept or lean toward internalism" and 28.9% answered "other." For discussion of the survey, see Bourget and Chalmers (2014). I think the issues are too complicated for either a crude externalism or a crude internalism to be true. However, the view I defend is most naturally classified as a kind of internalism.

water is H<sub>2</sub>O, people used “water” to talk about H<sub>2</sub>O while being ignorant of the fact that water is identical to H<sub>2</sub>O. Imagine that in 1750 Oscar, on Earth, believes that “water” is wet. Oscar has a duplicate on another planet, Twin Oscar, who (at the same time) also believes that “water” is wet. But on Twin Earth, the chemical structure of the stuff denoted by “water” is XYZ. So, the stuff Oscar talks about with “water” isn’t in the extension of Twin Oscar’s word “water,” and the stuff Twin Oscar talks about with “water” isn’t in the extension of Oscar’s word “water.”

Putnam also stipulated that Oscar and Twin Oscar were complete internal duplicates while different meanings attach to their tokens of “water”. Having stipulated that their ‘heads’ are the same while the meanings of their terms vary, Putnam concluded that linguistic meaning “ain’t in the head.” Unfortunately for the example, it’s not plausible that Oscar and Twin Oscar are truly duplicates if Oscar’s world includes H<sub>2</sub>O while Twin Oscar’s world includes only XYZ, for then Oscar and Twin Oscar are not really *physical* duplicates. Our bodies have lots of water (H<sub>2</sub>O) in them, so anything that is a duplicate of a normal human also has a body largely made up of water (H<sub>2</sub>O).

A better way to understand Putnam’s proposal is this: a speaker’s word “water” and concept *water* vary in extension depending in part on the environment in which the word and concept are acquired. So when Oscar and Twin Oscar both assent to the truth of the sentence “water is wet,” they assent to different *propositions*. Their *thoughts* are different. The thesis that implies is semantic externalism. Roughly, it holds that the content of a thought is constituted in part by entities in the thinker’s external

environment. This is the psychological corollary to the linguistic externalism Putnam initially advanced, which held that the meaning of a *term* is its reference.

It took a while before philosophers began to discuss the apparent fact that semantic externalism is incompatible with the idea that we can introspect to discover what we are thinking about. (It was 18 years from the publication of Putnam's 1973 to McKinsey's 1991.)<sup>13</sup> Consider the concept *water*. If the content of one's concept *water* is constituted in part by the environment one is in, then one cannot tell if one is entertaining the concept *water* without knowing something about the environment, namely, whether the environment includes water. So, one cannot examine one's concepts from the armchair. To put the problem a different way, if externalism is true and it is possible to examine one's concepts from the armchair, you can come to know that your environment contains water simply by introspecting on your concept *water*. That doesn't seem right.

So what is going on in Oscar and Twin Oscar's heads? More precisely, what mental contents correlating with their use of "water" should we assign to Oscar and Twin Oscar? A radical externalist reply would be while Oscar and Twin Oscar might assent to the same sentences involving "water," their *thoughts* invoking the concept *water* are utterly different. Putnam held that conclusion even though it is undermined by the very set-up of his story, as we just noted: Putnam claimed the twins were psychologically identical even though they must have been physically non-identical due to the presence of

---

<sup>13</sup> Externalists have various replies, at least some of which attempt to reconcile externalism with introspective access. See (McLaughlin 2003) for discussion. In the end, I think the original charge sticks.

H<sub>2</sub>O on one planet and its absence on the other. The consistent radical externalist story, then, is that the Oscars vary in their mental contents as a result of variations in their respective environments. Intuitively, though, Oscar and Twin Oscar *do* have something mental in common. Things seem exactly the same to them from the inside and they would assent to exactly the same descriptions of that which they call “water”. Therefore, radical externalism is false.

A reply on the very opposite end of the spectrum is radical internalism. On this view, Oscar and Twin Oscar have exactly the same mental content. Segal (2000) defends this view. Recall that Putnam’s thought experiment set the clock to 1750, prior to the discovery that our water is H<sub>2</sub>O. Plausibly, when we talk about water today we are talking about H<sub>2</sub>O. But, as Segal forcefully argues, it is far less clear that “water” referred *exclusively* to H<sub>2</sub>O prior to the discovery that the primary chemical structure of that substance that falls from the sky and fills the rivers is H<sub>2</sub>O.

To make his point, Segal advances a revised version of a thought experiment from Putnam (1975, p. 140). The year is 1750. An interplanetary shuttle service develops. It is discovered that both Earthlings and Twin-Earthlings use “water” to refer to the liquid stuff of their acquaintance. Both Earthlings and Twin-Earthlings are aware of their shared convention. As chemistry develops on each planet, it is discovered that the liquid stuff on Earth is made of H<sub>2</sub>O while the liquid stuff on Twin Earth is made of XYZ. Once that discovery is made, it is agreed by all that it would be too much hassle to make the switch to calling just one of these substances “water” and inventing a new word for the other, or inventing new words for both. So, it’s agreed that “water” refers to

whatever satisfies a description along these lines: the liquid stuff that falls from the sky, fills the rivers and lakes, and must be drunk for survival.

The point of Segal's story is that Putnam's story presupposed not only that "water" was a natural kind term prior to 1750 but also that "water" was a natural kind term that picked out *only* H<sub>2</sub>O. Segal's story suggests that "water" could be a natural kind term while allowing for a number of kind-properties to fall within the extension of the term. Linguistic meaning is, after all, a matter of convention. And there is nothing about "water" that *required* the community of experts to decide that "water" would forevermore be used to denote H<sub>2</sub>O rather than accepting a broader notion that could have included XYZ.<sup>14</sup> So, Segal suggests, in the actual world the meaning of water prior to 1750 *was not* uniquely H<sub>2</sub>O. The meaning was broader. Putnam's externalist conclusion, while possible, isn't required by imagining Oscar and Twin Oscar talking about "water" while the world each occupies has different substances playing the "water" role.

Segal goes on to argue that Twin Earth-type cases, while often compatible with externalist interpretations, are equally compatible with the radical internalist view that concepts are shared by Earthlings and Twin-Earthlings: they have the same extension conditions; at least they did prior to the narrowing of "water" to H<sub>2</sub>O here and XYZ there.

Radical internalism has, I think, a lot to recommend it. It may be true; I'm undecided. Fortunately, we need not decide for or against radical internalism. What is important for our purposes is that there be a way of making sense of the kind of content

---

<sup>14</sup> Searle (1983, p. 203) also makes this point.

Oscar and Twin Oscar *share* (on the assumption that “water” varies in extension between the two worlds) rather than arguing that the content they share is the only kind of content there is. The kind of content they share is called narrow content. The kind of content they don’t share—the kind of content we have in mind when we say “Twin Oscar isn’t *really* thinking of water because the watery stuff of his acquaintance is XYZ” is called ‘broad’ or ‘wide’ content.<sup>15</sup> A position in between radical externalism and radical internalism is a two-factor theory. Such a theory holds that there exist both narrow and broad content. What I’ll say about narrow content is compatible both with radical internalism and two-factor theory.

The possibility of armchair investigation into the nature of concepts requires that it is the narrow content of a concept that we investigate from the armchair. It should be clear enough that one cannot examine from the armchair the wide content of *water* thoughts: for such thoughts, if there indeed are any, have the content they do in virtue of features of the thinker’s environment. Surely one cannot discover facts about the environment external to one’s mind by introspection on non-relational features of one’s mind! While it may be true, in a sense, that Oscar and Twin Oscar entertain different propositions when they consider to themselves that ‘water is wet,’ the content of the concept of water each possesses has the same narrow content. The correct analysis of narrow content is controversial, as is how to understand the relationship between narrow

---

<sup>15</sup> The use of “broad,” “wide,” and “narrow” in the philosophy of mind literature strike me as especially poor terminological choices. There is no sense in which the difference in these types of contents is helpfully characterized as a difference in girth. Furthermore, philosophers differ over whether narrow contents exist, and over whether, if they do exist, they play a role in determining broad/wide content.

content and the propositions that serve as the objects of attitude constructions. I'll now briefly sketch what I think narrow content is, how it functions, and its role in determining the proposition expressed by a thought or sentence. Then we will return to the discussion of armchair conceptual analysis.

Recall that Putnam said Oscar and Twin Oscar were in the same psychological state even while their environments were different with respect to the presence/absence of water. Something would seem the same to Oscar and Twin Oscar about their thoughts of water. They use "water" the same way in sentences, entertain the same beliefs about "water", and so on.<sup>16</sup> This suggests that there is a layer of content shared by Oscar and Twin Oscar. Following Putnam, we'll call it narrow content. How to understand the nature of narrow content is controversial. But Putnam's examples provide us with a natural way of characterizing narrow content: it is a kind of descriptive content. Oscar's "water" thoughts are thoughts of *that liquid that flows from rivers, falls from the sky, is consumed to quench thirst, and that we, around here, call "water"*. Twin Oscar would associate with water the same description (read *de dicto*) that Oscar does.

We can say something similar about the content of definite descriptions. There is a clear sense in which the thought one has when one thinks, "The President of the United States is a very powerful person" is *the very same thought* whether the extension of "The President of the United States" is Barack Obama (as it was at the time I wrote this sentence) or whether it is someone else (as it will be two years from the time I wrote this

---

<sup>16</sup> They do not, according to the externalist, have the same beliefs about *water*, however: Putnam's very point is the at most one of them is thinking about water, even while both can truly be described as having beliefs about "water."

sentence). And one can *wonder* whether the President of the United States is a good person, *hope* that the President of the United States can deliver a strong economy, *believe* that the President of the United States should be more/less hawkish, and so on. We can easily make perfect sense of all of these attitude ascriptions even if we are ignorant of who is now the President of the United States. In this way, one can entertain *the same* descriptive thought in very different environments: e.g. an Obama ‘environment’ rather than an environment in which someone else is the President of the United States.

There is a complication here. The issue concerns the content of the description that serves as the narrow content of a thought. We just considered analyzing the narrow content of the concept ‘water’ as involving descriptions like: the *liquid* that flows in our *rivers*, that falls from the *clouds* in the *sky*, etc. An externalist might object that descriptive content does not adequately avoid external constituents: externalists who find plausible that “water” is a natural kind term that picks out H<sub>2</sub>O and only H<sub>2</sub>O are likely to find plausible the suggestion that the terms we used in our descriptions are equally natural kind terms whose meaning is determined by external factors. In other words, we haven’t given an adequate *internalist* descriptive analysis of “water” if our descriptions include concepts like “sky” because “sky” may have its meaning determined externally.

The solution, I think, is that the narrow contents of our thoughts consists in descriptions; and these descriptions, in turn, are composed of constituents with which the thinker is acquainted (Russell 1905, 1910). Each descriptor that appears in a descriptive thought must ultimately decompose into factors of which the speaker is *de re* aware. Russell thought he was acquainted with universals and sense-data, so the contents of

propositions he could understand reduce ultimately to those, plus the quantifiers and logical relations needed to complete the proposition.<sup>17</sup> We may or may not follow Russell with the *identification* of what it is thinkers are acquainted with: universals, sense-data, and his ‘logically proper names’ *this*, *that*, and *I*. The point here is that we can eliminate reference to allegedly external contents that appear in the description of a thought we can entertain by understanding the description in question to be one that is composed wholly of content constituted by entities with which we are acquainted.

While the correct analysis of propositions and thoughts is a matter of controversy, one way of translating the ideas expressed here is this. In the minds of speakers there are concepts. Concepts are constituted by factors in experience with which thinkers are acquainted. Thinkers entertain thoughts, which are constituted by descriptive contents that say something about how things are, and these descriptive contents are also constituted by factors with which the subject is acquainted.

From here, there is a choice: there is the radical internalist option and the two-factor option. The radical internalist will say thoughts determine extension: when a subject entertains a complete descriptive thought, that thought provides extension conditions. Take water, again. The radical internalist can hold that Oscar and Twin Oscar think exactly the same thing, so there is no difference in what Oscar thinks and what Twin Oscar thinks. It would be false to say “Oscar is thinking about water, but Twin Oscar isn’t, because ‘water’ picks out XYZ in his environment, and XYZ isn’t

---

<sup>17</sup> I am using ‘proposition’ in the broad way I have used it earlier in this section: Russell’s view of propositions was unique and changed substantially over time.

water.” For each Oscar, the thought he entertains is identical with the proposition that describes his mental state. They do not come apart.

Alternatively, there is the two-dimensionalist option: one and the same descriptive *thought* expresses different *propositions* in different environments.<sup>18</sup> Oscar and Twin Oscar think the same thought, but their thoughts express different propositions. The difference between these views hinges ultimately upon whether we want to follow the externalist in holding that propositions ever involve mind-independent entities as constituents. The reason the *propositions* expressed by Oscar and Twin Oscar’s type-identical thought differ is that the descriptive thought picks out H<sub>2</sub>O for Oscar and XYZ for Twin Oscar. One way to describe the relation between the propositions is this. Each Oscar entertains a thought. But his thought is indeterminate between a number of propositions, depending on what his narrow concept *water* picks out in his environment. Factors external to the minds of each, plus the thought each Oscar thinks, determine a proposition. Propositions, on this view, are not identical with thoughts, and thinkers uttering meaningful sentences will often not have a determinate proposition in mind when uttering a sentence.<sup>19</sup>

---

<sup>18</sup> I take my description of the two-factor view to be structurally very similar to the kind of two-dimensionalist views defended by (Chalmers 1996, 2003) and (Jackson 1998), among others. The main difference is that I am distinguishing ‘thought’ and ‘proposition’ while the other views distinguish two kinds of proposition; nothing hinges on the choice of one terminology over another.

<sup>19</sup> Two-dimensionalism says two different propositions are expressed by thought ascriptions like “Oscar is thinking that water is wet.” That is possible only if the sentence has two semantic values giving rise to two different propositions being expressed in any context of utterance. But the sentence does not appear to have two semantic values. I don’t think that is decisive against two-dimensionalism. But it does count in favor of radical internalism.

I am trying to remain neutral about whether to understand propositions in such a way that speakers are often aware of some semantic value of the sentences they assert but not aware of the propositions they seek to entertain, or the radical internalist way, where the content of thought is exhausted by internal factors and so there is no distinction to be drawn between thoughts and propositions. It doesn't matter for my purposes. All I need is to isolate a layer of content that speakers can access and understand. Philosophers can only examine from the armchair whatever occupies that 'layer.'

The purpose of the discussion in this section up to now has been to address the objection from semantic externalism that we cannot analyze our concepts from the armchair because our concepts are often constituted by mind-independent factors, therefore we have to get outside of the mind to analyze the extensions of our concepts. Thus Putnam: "Meanings ain't in the head!" I've attempted, first, to undermine one of the most compelling thought experiments for externalism; second, to say what narrow content is supposed to be—namely, descriptive content; third, to consider what options there are for the defender of narrow content—there are two, radical internalism and two-factor theory. Now to wrap it up by answering the initial concern: how does all this allow us to analyze our concepts from the armchair?

We have already identified the conceptual content that can be analyzed from the armchair. That is narrow content. But I said above that we often—perhaps *usually*—examine our concepts largely by means of considering what we would say using certain terms embedded in sentences that are used to describe a hypothetical state of affairs. What is the connection between narrow concepts and words that allows this to work?

All I can do here is very briefly sketch a just-so story of how concepts come to have contents and words come to express concepts. Individual minds are acquainted with various phenomenal properties: shapes, sounds, feels, what-its-likenesses. Minds classify the characteristics of their experiences. These classifications are individualistic. It would be extremely difficult to fully verbalize the richness of one's conscious experience. But we can attempt to articulate the character of our conscious states by making up words to stand for classifications of conscious experiences. Others have experiences much like ours: experiences they can recognize as roughly fitting the descriptions we offer of our experiences. Words in a public language acquire meaning by being used to pick out these shared features of conscious states the richness of which outstrips the content of the concept expressed by the word.

There are four 'layers' in this story:

(1) *Phenomenal experience*: minds have experiences.

(2) *Classifying experiences into concepts*: minds classify these experiences, using concepts to pick out common features. Some of the features classified have an essentially phenomenal character: e.g. the concept I would describe as *this red*, where all I can do to classify the experience is point to a particular instance familiar to me. However, other concepts, though still private, do not have such a robust phenomenal character: the concept *I* seems to be of this latter sort.

(3) *From private concepts to public concepts*: our concepts are shaped by our interactions with other people. As a result of such feedback, we settle on certain contents as 'the' concept of such-and-such. For example, I may have learned the concept *apple* as

a result of my parents' pointing to certain kinds of apples—perhaps all we ever had at my house were red delicious apples—but through interactions with other people I learn that the concept *apple* has a more general character than just the instances I had encountered.

(4) *Public language*: many of our concepts come to be expressed by words in public language. The words we use to express our concepts often fail to express precisely the character we associate with those concepts. If, in fact, my concept *apple* did include an association with only red fruit, that content would not be expressed by my use of the word “apple” when I speak with others. The content that is expressed by the use of a word in language is the result of a kind of miracle of communication: speakers intend to express a content and manage to be understood by others; repeated application of this process results in public meaning.

This way of thinking about how thoughts come to have content, and about what those contents are, is an attempt to briefly sketch a Gricean view about meaning (Grice 1957). On that view, the meaning of a word or sentence is determined by speakers with communicative intentions attempting to express their thoughts to others with the goal of getting others to recognize what they (the speakers) *intend* to express through their words and sentences. The meaning of a public-language sentence reduces to the meaning individual speakers attempt to express through their uses of language.

However, it doesn't follow that we always successfully express with language what we intend to express. The thought one thinks is constituted by the concepts one employs in the thought. When I think the thought I would describe as the thought *that Oscar is thinking about water* the constituents include my concept of Oscar and my

concept of water. That has no direct bearing on what is asserted by me if I say, “Oscar is thinking about water.” If my concept of water were idiosyncratic, it would have no effect on what proposition is expressed by a *sentence* I utter that includes the word “water”. Just as one’s concepts can fail to capture the richness of phenomenal experience, one’s sentences can fail to express one’s thoughts. That will not affect what one’s linguistic utterances mean because what words mean is determined by one’s language-using community. One cannot infallibly express one’s thoughts by using language however one wishes.

Given that there is some loss of content between levels, it is not surprising that we cannot simply introspect the richness of our higher (3) concepts, whether those of philosophical interest, like *knowledge*, or the commonplace, like *sandwich*. This account explains why, as I have emphasized, it is difficult to provide an analysis of concepts of philosophical interest. Now I will explain why competent speakers are able to make serious gains in providing an analysis by reflecting upon what they are inclined to say about hypothetical scenarios.

The ability to communicate requires that speakers in the same community using the same language agree about what concepts terms in their language express and what thoughts are expressed by grammatically complete expressions in their language. The more commonly used a term is, the more expert speakers will likely be about precisely what concept is expressed by a term. With a language as sophisticated as English, our terms can be used to express very complex thoughts and concepts. But in every case, because language-users are applying concepts constructed out of experience, the content

of their concepts is within their reach, so to speak. So, speakers can be *competent* speakers partly because meanings are constructed out of factors in phenomenal experience. That explains why competent speakers are in a position to analyze their own concepts.<sup>20</sup>

The method we can employ to analyze our own concepts is essentially Carnap's. Carnap (1955) argued that we can understand a speaker's concept C (in his terminology, 'intension') by considering all of the things of which the speaker is disposed to predicate C. So far, I have been emphasizing that we can do this sort of thing with our own concepts, where the speaker in question is oneself. *I* can consider all of the things *I* am disposed to predicate C of and thereby come to propose an analysis of C. And, on the very plausible assumption that I am a normal competent speaker—meanings are in my head, and I seem to get by just fine communicating with others—and that I am using terms to express the same concepts as others in my language community, I will, by looking inward, have discovered the correct analysis of the concept C as understood by my community. That is, I will have discovered the analysis of C.<sup>21</sup>

---

<sup>20</sup> In a position, that is, as far as the concepts are concerned. They may be uninterested in doing so, intellectually incapable, or otherwise not in a position to do so for a host of reasons.

I argue in chapter four that my conditions for semantic competence are compatible with semantic externalism, strictly speaking. While that is the case—*strictly speaking*—externalists owe an account of why we should think speakers are generally competent with language given that they grant that meanings are often cognitively inaccessible. *Prima facie*, externalism seems to make speaker competence mysterious. Cf. (Kripke 1980, esp. Lecture III, Putnam 1975).

<sup>21</sup> I talk of analyzing *the* concept such-and-such in the footnoted sentence and elsewhere in this dissertation. It is likely that, for most Cs, my concept C is a little different from the concept C others have now, or the concept C that I had a month ago, and a year ago, etc. Nevertheless, I will assume that I am a normal speaker. This assumption is not a shot in the dark, either. As I mentioned above, I seem to get along just fine communicating with others, my students mostly follow my lectures, my peers usually find my thought experiments at least a little plausible, etc. So, by figuring out the contents of my own concepts, it is reasonable to assume that I will have figured out something significant about the way *we* think.

#### 1.3.4 *The Analytic/Synthetic Distinction, A Priori Justification, and Necessity*

As I said earlier in the chapter, one might understand “giving an analysis” in two different ways. One is as the offering of analytically necessary and sufficient conditions for some concept. The other is as an attempt to understand the concept in question, with no claim to have offered analytically necessary sufficient conditions. If the first route is at all viable, then there has to be such a thing as the analytic/synthetic distinction. After all, there are many claims that can be made using a concept, and they seem to differ in status. “Bachelors are unmarried” seems to have a different status than “bachelors are lazy,” and many philosophers have thought the difference in status owes to a special property possessed by “bachelors are unmarried” but not possessed by “bachelors are lazy.” If someone says that is because the former is analytic while the latter is not, then there had better be some principled way of spelling out that distinction. And it is well known among philosophers that the distinction has been challenged, especially by Quine.

I think it is acceptable to document a search for an analysis that fails to discover analytically necessary and sufficient conditions for some concept without having to declare the endeavor an overall failure. The search can be illuminating nonetheless. The success of this dissertation does not depend on defending analyticity. One can surely offer conditions for knowledge that say when one has knowledge and when one doesn't without claiming the proffered conditions are analytic. However, I also think the analytic/synthetic distinction is intuitive and defensible. Now I'll say why. Along the way, I'll discuss the closely related notions of necessity and *a priori* justification.

#### 1.3.4.1 *The Analytic/Synthetic Distinction and the A Priori*

Much of the criticism of the analytic/synthetic distinction is, regrettably, extremely confused, often the result of applying criteria that are simply not appropriate for arguing for or against the existence of the distinction. For a sample, consider the opening paragraph of the article, “The Analytic/Synthetic Distinction,” in the Stanford Encyclopedia of Philosophy.

“Analytic” sentences, such as “Ophthalmologists are doctors,” *are those whose truth seems to be knowable by knowing the meanings of the constituent words alone*, unlike the more usual “synthetic” ones, such as “Ophthalmologists are rich,” *whose truth is knowable by both knowing the meaning of the words and something about the world*. Beginning with Frege, many philosophers hoped to show that knowledge of logic and mathematics and other apparently *a priori* domains, such as much of philosophy and the foundations of science, *could be shown to be analytic by careful “conceptual analysis.”* This project encountered a number of problems that have seemed so intractable as to lead some philosophers, particularly Quine, to doubt the reality of the distinction. Surprisingly, this led him and others to doubt the reality and determinacy of psychological states. There have been a number of interesting reactions to this scepticism, both in philosophy and in linguistics, but, while the reality of psychological states might be saved, it has yet to be shown that the distinction will ever be able to *ground the a priori in the way that philosophers had hoped*. (Rey 2013, italics mine)

The three italicized bits deserve comment, in order. First, analyticity is here described in *epistemological* terms: analytic sentences are ones the truth of which one comes to *know* a certain way, while the truth of synthetic sentences is known a different way. (It should say “the truth *or falsity*” is known a certain way: e.g. unicorns have two horns.) Second, it is mentioned that some philosophers are concerned to *demonstrate* that some sentences are analytic. Third, the problem is raised that some are skeptical that analyticity can

“ground the *a priori*,” which seems to be something analytic truths might be thought useful for, but is clearly not constitutive of analyticity itself.

The main problem here is that a metaphysical matter is repeatedly cast as an epistemological one. Whether or not a sentence is analytic is best understood as a property of the sentence, presumably an intrinsic property rather than a relational one the possession of which depends upon minds knowing that the sentence has some other property or properties. The question of how minds are related to analytic sentences—that is, how anyone could come to know or justifiably believe that a sentence is analytic—is a separate matter from what analyticity *is*. *A priori* might be the sole *criterion* of analyticity, but that does not make *a priori* identical to analyticity. Now, one might want to argue that a sentence is analytic if and only if it is *a priori*, where the fact that the sentence has the relational property *being knowable a priori* is taken to guarantee its having the intrinsic property *being analytic*.<sup>22</sup> I’ll return to this shortly.

To address this confusion requires attention to the nature of analyticity. After doing so we will be in a position to address the epistemological status of analytic sentences.

One must separate the analytic/synthetic distinction from the various other properties philosophers have often thought are closely connected with analyticity. The clearest statement of the distinction is from Frege (1884 (1997), sec. 5). Frege held that a sentence is analytic when it is reducible to a truth of logic by substituting synonyms for

---

<sup>22</sup> Strictly, the quote from Rey that opens this section is compatible with a ‘criterion’ reading rather than the ‘analysis’ reading that I criticize here. However, if that is so, he gives no indication anywhere in the article that he intends the criterion reading over the analysis reading.

synonyms. For example, “all bachelors are unmarried” is analytic because bachelor means “unmarried male eligible for marriage”; so, substituting ‘bachelor’ in “all bachelors are unmarried” yields a logical truth: “all *unmarried* males eligible for marriage are *unmarried*.”<sup>23</sup> That is equivalent to the tautology:

$$\forall x((Ux \& Ex) \rightarrow (Ux))$$

A sentence is synthetic when it is not analytic.

That’s the distinction. It is plausible to think that many sentences are indeed analytic. Identifying philosophically interesting examples is tricky business, however. The reason for this difficulty is straightforward. In the “all bachelors are unmarried” example, we assume that ‘bachelor’ is *synonymous* with “unmarried male eligible for marriage.” For that explanation to work, ‘unmarried’ must be part of the meaning of ‘bachelor’. But what could it mean to say that *part of the meaning* of one expression is some other meaning?

Something like Kant’s much-derided early idea of analytic sentences as those where the concept of a predicate is ‘contained in’ the concept of the subject is, I believe, the answer to that question (1781 (2006), §A7/B10). The containment account offers a straightforward answer to the question I just asked, namely, How could part of the meaning of one expression be some other meaning? As helpful as Kant is on that point, the account has a well-known defect. Unfortunately, one of the problems with Kant’s

---

<sup>23</sup> I won’t argue for this analysis of ‘bachelor’ beyond expressing dissatisfaction with the simple ‘unmarried male’ analysis of bachelor. If a contestant on a dating game were told that the bachelor behind curtain 1 is fun, single, and athletic, loves to read and cuddle, and has a great laugh, and after lowering curtain 1 she were to find my three-year-old son sitting there, she would justifiably be miffed. Plausibly, my son, though an unmarried male, is not a bachelor. Neither is the Pope, for a similar reason.

account of analyticity is that it fails to identify many intuitively analytic sentences as analytic. His account works for “all bachelors are unmarried,” because plausibly ‘unmarried’ is part of the meaning of bachelor. But his account struggles with lots of other examples that are analytic according to Frege’s account of the distinction. Frege said a sentence is analytic when it is equivalent to a logical tautology upon substituting synonyms for synonyms. That definition is incompatible with Kant’s account because Kant’s account is limited to sentences with a straightforward subject-predicate structure. With a little creativity, it is easy to identify Frege-analytic sentences that are not Kant-analytic: “there are no round squares,” “Hitler was a Nazi or (it is raining or it is not raining),” “unicorns have horns and (if unicorns do not have horns then  $2+2=4$ ),” etc. All of these examples defy the subject-predicate sentence structure Kant was preoccupied with.

Nevertheless, we can make good use of Kant’s insight. Philosophers searching for an analysis of some concept want to locate true identity statements that explicate concepts that have captured their interest. We want to find what is to be substituted for  $x$  in sentences with this form: “knowledge is  $x$ ,” “good is  $x$ ,” “free action is  $x$ ,” and so on. We assume that what goes for  $x$  is *not* itself a tautology: while ‘ $y$  is  $x$ ’ is necessary, ‘ $x$ ’ *by itself* is not. The containment account is, as far as I can see, indispensable for understanding how ‘ $y$  is  $x$ ’ can be analytic while ‘ $x$ ’ isn’t: it is because the concept expressed by ‘ $x$ ’ is *contained in* ‘ $y$ ’. So, ‘ $y$  is  $x$ ’ *means* (something like) ‘ $x+a+b$  is  $x$ ’. Something that appears on the left-hand side of the identity sign appears again on the right-hand side. But that would be impossible unless the analyzed complex concept

*contains* a simple constituent that is repeated on both sides of the identity sign: on one side as part of the analyzed complex, and on the other as an isolated part.

That this is so explains how ‘substituting synonyms for synonyms’ works for sentences that are not superficially tautologous (like “an unmarried man eligible for marriage is an unmarried man”). If “all bachelors are unmarried” is analytic because the analysis of “bachelor” makes ineliminable reference to “unmarried,” *part of the meaning* of ‘bachelor’ has to be ‘unmarried.’ And if part of the meaning of ‘bachelor’ is ‘unmarried’, the concept *unmarried* is contained in the concept *bachelor*. While many analytic sentences are not susceptible to analysis in terms of the containment metaphor, analytic sentences that are philosophically interesting will be. We want to identify philosophically interesting analytic identity statements, and those must be understood in terms of conceptual containment.

Understood as a claim about structural constitution, the concept containment claim sounds downright bizarre. No ontological view about the nature of concepts will permit that one concept can be literally inside of another concept, at least not in the same way that a person sits inside of a car. But surely the containment claim should not be understood in this way. The containment claim is not problematic if we understand ‘containment’ to be a matter of the *content* of a concept. Part of the content of the concept ‘bachelor’ is the concept ‘unmarried.’ Plausibly, the containment idea does require, however, that complex concepts possess structure, so the meaning of ‘unmarried’ is a *constituent* of the meaning of ‘bachelor,’ but not in virtue of being inside of ‘bachelor.’

Having stated the analytic/synthetic distinction, let us move to a more controversial issue: the epistemology of analytic sentences. Is it the case that all analytic sentences are *a priori* and that all *a priori* sentences are analytic? We can begin to answer that question by considering again “all bachelors are unmarried.” The meaning of ‘bachelor’ is an *a posteriori* discovery: one comes to discover the meaning of that term by interacting with a linguistic community. But once that concept is acquired, plus the concepts *all*, and *are (is)*, there seems to be a distinct way to come to justifiably believe “all bachelors are unmarried,” a way seemingly unlike the way one comes to hold beliefs about the weather by walking outside. Rather, by simply considering the statement “all bachelors are unmarried,” one comes to have an extremely good reason for believing that sentence is true.

There is an argument that the kind of ‘striking one as true’ that occurs when one reflects, perhaps for the first time, on simple analytic sentences like “all bachelors are unmarried” cannot produce any *new* epistemic justification for believing the sentence is true. The argument is due to Gillian Russell, and it goes like this.<sup>24</sup> To understand the sentence “all bachelors are unmarried” one must possess all the concepts expressed in that sentence. But if that sentence really does reduce to a truth of logic as the Fregean criterion says it must if it is analytic, then the concept *unmarried* is already expressed by ‘bachelor.’ As a result of your competence with *bachelor*, you have no new reason to

---

<sup>24</sup> The argument as I present it is due to Gillian Russell, though she has distinctive externalist commitments underlying her view of analyticity that I reject (Russell 2008).

believe “all bachelors are unmarried” is true upon considering the sentence than you did when you first acquired the concepts *all*, *are (is)*, *bachelor*.<sup>25</sup>

The argument’s conclusion appears to generalize. Given that competence with a concept that includes as part of its content the content of another concept is necessary to understand an analytic sentence, consideration of an analytic sentence cannot produce any new justification for the person considering the sentence. Therefore, reflection on analytic sentences cannot produce *a priori* justification.

While the argument says something very interesting and, I think, true about the nature of philosophical analysis, the anti-apriorist conclusion does not follow. A position more friendly to the *a priori* that escapes the argument is that *a priori* justification that relies essentially on the consideration of analytic sentences is always trivial in nature, in the sense that one gains no new *information* about one’s concepts when considering analytic sentences. Instead, one comes to appreciate in a new way the manner in which concepts one already possesses are related. So, one could have better justification for believing “all bachelors are unmarried” upon considering the sentence and having the experience of being struck with its truth than one did before considering the sentence but while possessing the concepts. Metaphorically, one ‘puts together’ that a man is a bachelor only if he is unmarried, even if one has never explicitly entertained the thought before. I can own all the lumber, screws, glue, tools, and know-how necessary and

---

<sup>25</sup> Arguably, Russell’s concern is a version of the paradox of analysis, aimed specifically at the allegedly *a priori* status of philosophical analyses. The paradox of analysis asks how philosophical analyses can strike us as informative if philosophical analysis consists in unpacking concepts we are already familiar with. As a result, philosophical analyses should be either obvious or false.

sufficient to build a table, but until I build the table, I don't have the table. I can possess all the concepts necessary and sufficient to understand an analytic statement, but until I put all that together and consider the sentence, I lack *a priori* justification for believing the sentence.

That brings us to the truth in (Gillian) Russell's argument. My reply to the argument presupposed that one's concepts and their relations to one another—even relations that obtain as a matter of the constitution of the concepts in question—can be at least partly hidden from a person who is a competent user of terms in the language that express those concepts. Earlier I characterized competence with a concept in a dispositional way: one understands a term when one *would* have certain reactions in certain circumstances. With such conditions for concept possession it should come as no surprise that one would have the disposition to immediately assent to "all bachelors are unmarried" even while having competence with 'bachelor' and 'unmarried' and while having never explicitly considered "all bachelors are unmarried." On the view of the *a priori* I am currently advancing, a discovery can be *a priori* even though one gains no new information about the truth of the analytic sentence, and even though the 'discovery' consists in learning that a logically trivial sentence is true.

Recall the Fregean criterion, which says a sentence is analytic if and only if it reduces to a truth of logic by substituting synonyms for synonyms. Once synonyms are substituted for synonyms, one is left with a formal tautology. If that tautology is not *a priori*, then the analytic statement that was reduced to a tautology won't be, either. So, the logical basis of the *a priori* justification of analytic sentences is in the logical

triviality that underlies analytic statements. Given this unsurprising result, it is no objection to the view sketched above that it makes the *a priori* of analytic statements logically uninformative. Of course, the fact that analytic sentences are *logically* uninformative is no barrier to their being *epistemically* very informative. There are all sorts of logically necessary truths that are very difficult to understand, let alone justifiably believe. But because being *a priori* is distinct from being analytic, it is easy to make sense of a sentence being logically trivial even while coming to justifiably believe that it is true is a serious achievement.<sup>26</sup>

Plausible examples suggest that the domain of *a priori* justification is not restricted to analytic sentences. It seems just obvious that the following is *a priori*: “nothing is both red all over and green all over (at the same time and in the same sense of ‘all over’).” But attempts to reduce that statement to a tautology fail. It seems there is little left to do than to simply admit that the statement is both synthetic and *a priori*. Interestingly, the same account of concept possession and language understanding I have offered not only allows for *a priori* justification of synthetic statements, but also permits that it, like analytic *a priori* justification, does not provide the person who acquires that justification any new information about the concepts related in the synthetic sentence.

---

<sup>26</sup> Laurence Bonjour, in his *In Defense of Pure Reason* (1998) is criticizes a view he calls Moderate Empiricism. “[M]oderate Empiricism can be understood as an attempt to defend two main theses...first, *a priori* justification is limited to analytic propositions...second...[it] does not require the sort of allegedly mysterious intuitive capacity advocated by rationalism...” (29). Overall, I think Bonjour’s arguments that *a priori* insight is necessary for *a priori* justification are successful. But he ignores the plausibility of a position according to which the insight in question is identical to understanding language and formal tautologies on his way to endorsing a kind of Platonism that is, I think, superfluous. So, Bonjour is right that just saying ‘*a priori* justification only applies to analytic statements’ is epistemologically inadequate, but it doesn’t follow that the most plausible version of moderate empiricism is thereby refuted. In the following remarks in the text I extend this idea to the synthetic *a priori*.

In the sense of ‘new information’ we have used in the discussion of the *a priori* thus far, one who considers that statement for the first time acquires no new information about the way red and green are related to one another. The subject who understands the statement is competent with both *red* and *green* and simply recognizes upon considering the statement that the statement has to be true in virtue of the way *red* and *green* relate to one another, even though that relation is not one of partial identity through concept containment. It is, again, understanding concepts that provides the basis both for the truth of the judgment and for the epistemology of realizing *a priori* that the statement is true. So, we recognize inclusion and exclusion relations between concepts that are necessary, but non-formal: that is, not reducible to formal tautologies.

Attention to the fact that our concepts are often partly obscured from direct inspection not only makes sense of the phenomenon of recognizing *a priori* that a sentence is true, but also makes sense of the confusions uncovered in the quotation at the opening of this section, where the concepts of *a priority*, analyticity, and obviousness were woven together. The fact that a sentence is analytic, as we saw, has little to do with what it takes to come to justifiably believe that the sentence is true. As a result, the fact of disagreement about the *a priority* of a sentence is *far* from decisive against a sentence’s putative analyticity. The fact of disagreement over the analyticity of a statement is also far from decisive against the analyticity of a statement. It takes extremely careful reflection and, often enough, argument to come to discover that a statement is analytic. That’s the solution to the paradox of analysis I favor: analytic

sentences are often in a sense uninformative, but they are not epistemically trivial. (Cf. Fumerton (1983).)

I won't defend a position on the nature of *a priori* justification here. My concern has only been to spell out how the account of philosophical analysis described thus far relates to traditional concerns about the *a priori* nature of philosophy. I have tried to show two things here. First, that the relationship between philosophical analysis and concerns about the *a priori* are complicated; second, that offering a philosophical thesis as an analysis in the strict sense of an allegedly analytically true statement is defensible.

#### 1.3.4.2 *Armchair Philosophy and Necessity*

Since the rise of semantic externalism a new problem has been posed for the basic method of thought experiment. A version of the problem affects my semantic internalist view, too, which is part of the reason this problem is being raised after a discussion of the *a priori* and philosophical analysis.

The problem for externalists concerns the relationship between conceivability and possibility. The traditional view of the relationship is perhaps most evident in Descartes, who thought the fact that he could conceive of mind as separate from body entailed that it is possible for mind and body to be distinct; and if it is possible for mind and body to be distinct, then embodiment is not essential to mind; therefore mind and body are actually distinct.

[O]n the one hand I have a clear and distinct idea of myself, in so far as I am simply a thinking, non-extended thing [that is, a mind], and on the other hand I have a distinct idea of body, in so far as this is simply an extended, non-thinking thing. And accordingly, it is certain that I am really distinct from my body, and can exist without it (AT VII 78: CSM II 54).

A view like Descartes' presupposes that the natures of mind and body are open to direct inspection by the philosopher. It is by first challenging this assumption that the semantic externalist begins to sever the connection between conceivability and possibility. Recall that Putnam holds that water is identical to H<sub>2</sub>O. It follows that it is not possible that water is not H<sub>2</sub>O, regardless of whatever anyone conceives or imagines about the possible constitution of water. Even though it is conceivable that water is not H<sub>2</sub>O, it is not possible that water is not H<sub>2</sub>O. As Putnam baldly put it, "It is conceivable but it isn't logically possible! Conceivability is no proof of logical possibility" (Putnam 1975, p.151).

Saul Kripke called identities like the identity of water with H<sub>2</sub>O 'a posteriori necessities': the statement "water is H<sub>2</sub>O" is necessary because water is *essentially* H<sub>2</sub>O, and the relation is not discoverable by reflection on concepts, and thus it is a posteriori (Kripke 1980). Kripke dubbed 'metaphysical necessity' the kind of necessity that connected *being water* with *being H<sub>2</sub>O*. Kripke's 'metaphysical necessity' and 'metaphysical possibility' is what some other philosophers, including Putnam, called 'logical necessity' and 'logical possibility.'

Kripke devotes some effort in *Naming and Necessity* to first explaining why metaphysical necessities generally cannot be discovered by attempts to conceive of various scenarios in which, for example, water is not H<sub>2</sub>O; and second, to attempting to recover the connection between conceivability and possibility for some metaphysically necessary connections: notably, in an argument from the conceivability of mind-body dualism to the actuality of mind-body dualism.

I won't discuss the merits of Kripke's attempt to salvage the conceivability-possibility link for certain claims of metaphysical necessity, nor similar attempts by more recent authors.<sup>27</sup> It should be clear enough why the problem Kripke raises needs a response, however: if the things falling in the extension of one's concepts can have properties that are not discoverable by the possessor of those concepts, then in principle one cannot infer from the fact that one's concepts are (or are not) related in a certain way to the claim that the things in those concepts' extension are (or are not) related in that way.

The semantic externalists' version of this problem is particularly drastic because externalists think demonstrative and descriptive thoughts pick out kinds of entities directly. Putnam thought that prior to the discovery that the molecular constitution of water is H<sub>2</sub>O whenever anyone said (pointing to water) "*that* stuff," or made reference to water by means of a descriptive expression—"the liquid in the river is cold"—the speaker picked out the natural kind H<sub>2</sub>O. The same considerations, externalists hold, allow speakers to pick out other kinds. Perhaps "mind," "body," "knowledge," are all natural kind terms. In general, when our concepts pick out natural kinds, the inference from conceivability to possibility is dubious.<sup>28</sup> We cannot, when using kind-concepts,

---

<sup>27</sup> Cf. especially several of the papers in Gendler and Hawthorne (2002).

<sup>28</sup> This is similar to the modal objection to identifying the meaning of a name (e.g. "Aristotle") with a definite description (e.g. "the student of Plato") Kripke discussed in *Naming and Necessity*. That objection is that claims involving names have different modal properties than claims involving allegedly identical definite description. But that shouldn't be possible if the names are equivalent in meaning to the description. E.g. possibly, Aristotle was not the student of Plato; but it is not possible that Aristotle was not Aristotle.

run our thought experiments and construct philosophical analyses without detailed scientific knowledge.

In a previous section I challenged some arguments for semantic externalism. Here I want to do something a little different: challenge the idea that when we analyze our concepts we are thinking of any sort of necessity other than how our concepts are related. An assumption underlying Kripke-Putnam externalism is that there are necessities in nature and that some—indeed, many—of our terms refer to entities standing in relations that obtain necessarily. Recall the thought experiment Segal suggested as a response to Putnam’s story of discovering that “water” is used by inhabitants of Earth to refer to H<sub>2</sub>O while it is used by Twin Earthers to refer to XYZ. Segal suggested that the question of what semantic policy to adopt—for everyone, Earthers and Twin Earthers alike, to use “water” only for H<sub>2</sub>O, only for XYZ, or for both H<sub>2</sub>O and XYZ—is just a matter of convention. That idea ultimately supports not only the idea that our terms *don’t* pick out natural kinds directly, but also that the sort of necessary connection that might obtain between water and anything else is a matter of how we define “water”. The same goes for other allegedly natural kind terms, names, demonstratives, and rigidly-designating descriptions.

If some scenario like the one Segal imagines were to obtain, it would be an open question how we are to define “water” going forward. Similarly, Putnam has asked us to imagine that we could discover that all the things we previously called “cats” are not animals, but cleverly disguised robots (1962). Were we to learn such a thing, we would have a choice: either to understand “cat” so it continues to apply to those creatures, so it

will be true that cats are robots; or stick with ‘cats are animals’ and conclude that there are no cats.

The fact that we have this sort of choice is significant. It suggests that however things are with mind-independent reality, there is always a prior *semantic* question of how to categorize things. The truth or falsity of “necessarily, all cats are animals” depends upon our concepts of *cat* and *animal*, regardless of what relations obtain between properties in the world. Before getting worried about how to discover *a posteriori* necessities like ‘water=H<sub>2</sub>O’, we should pause to notice that the truth of that identity statement depends upon what kind of mind-independent modal connections we think can obtain between properties.

One of the assumptions of Kripke-Putnam externalism is that there are necessary connections that obtain between natural kinds. Hume famously challenged the notion that we have any idea of necessary connections that obtain independently of thought. I’m not going to outright endorse the Humean view. But Hume’s view of necessity allows for a clear and principled connection between conceivability and possibility, a connection that is rather difficult to make sense of on the assumption of metaphysical essentialism that is key to semantic externalism.

Here are a few passages from Hume about our understanding of necessity.

‘Tis an establish’d maxim in metaphysics, *That whatever the mind clearly conceives includes the idea of possible existence*, or in other words, *that nothing we imagine is absolutely impossible*. We can form the idea of a golden mountain, and from thence conclude that such a mountain may actually exist. We can form no idea of a mountain without a valley, and therefore regard it as impossible. (Hume 1739 (2000), I.2.32)

Our foregoing method of reasoning will easily convince us, that there can be no *demonstrative* arguments to prove, *that those instances, of which we have had no experience, resemble those, of which we have had experience.* We can at least conceive a change in the course of nature; which sufficiently proves, that such a change is not absolutely impossible. To form a clear idea of any thing, is an undeniable argument for its possibility, and is alone a refutation of any pretended demonstration against it. (Hume 1739 (2000), I.3.6.5)

I am, indeed, ready to allow, that there may be several qualities both in material and immaterial objects, with which we are utterly unacquainted; and if we please to call these *power* or *efficacy*, 'twill be of little consequence to the world. But when, instead of meaning these unknown qualities, we make the terms of power and efficacy signify something, of which we have a clear idea, and which is incompatible with those objects, to which we apply it, obscurity and error begin then to take place, and we are led astray by a false philosophy. (Hume 1739 (2000), I.3.14)<sup>29</sup>

All the objects of human reason or enquiry may naturally be divided into two kinds, to wit, *Relations of Ideas*, and *Matters of Fact*. Of the first kind are the sciences of Geometry, Algebra, and Arithmetic; and in short, every affirmation, which is either intuitively or demonstratively certain... Though there never were a circle or triangle in nature, the truths, demonstrated by EUCLID, would for ever retain their certainty and evidence.

Matters of fact, which are the second objects of human reason, are not ascertained in the same manner; nor is our evidence of their truth, however great, of a like nature with the foregoing. The contrary of every matter of fact is still possible; because it can never imply a contradiction. (Hume 1777, IV.1)

In these passages Hume says a few important things that are relevant to my discussion.

First, he claims that whatever we can conceive is possible and that whatever we cannot

---

<sup>29</sup> There is a tradition in Hume interpretation according to which Hume accepted the existence of 'secret powers' in nature that govern regularities we experience: cf. (Wright 1983b) for a defense of the view. I confess I find the view borderline incoherent. According to it, Hume undermines any reason we could have to think there are necessary connections in nature, and then goes on to posit them. If the view just said Hume is silent about what 'secret powers' there may be in nature, that would be fine. After all, Hume's argument is that we have no basis for saying what necessary connections there may be in nature one way or the other.

conceive is impossible. Second, he applies the conceivability test to “the course of nature”, where he claims it is possible that the natural world could change on account of the conceivability of such a thing happening. Third, he allows that there may exist so-called natural powers, but argues that if there are any such things, we have no idea what they are. Fourth, he distinguishes matters of fact from relations of ideas, and claims that the contrary of every claim of fact is possible.

It is clear that Hume believed there were necessary connections and that they could be investigated. It is consistent with such a view that ‘water=H<sub>2</sub>O’ could have been necessary: that is, if and only if it were inconceivable that something could be water without being H<sub>2</sub>O, then it would have been necessary that water is H<sub>2</sub>O. However, the essentialist denies that conceivability generally reveals necessary connections in this way. Essentialists like Kripke and Putnam claim that there are necessary connections that often fail our conceivability test, but which are discovered by science.<sup>30</sup>

Hume and contemporary essentialist externalists are obviously opposed on the issue of whether or not there are necessary connections between natural kinds that are not discoverable by the conceivability test. As Hume claimed above, there may be connections between properties in the world that we are unaware of—but whatever they might be, we can say nothing about them. It’s important to note on this point that essentialist externalists want something stronger: they claim there are necessary connections not discoverable by the conceivability test that we can discover by other

---

<sup>30</sup> An important question for an essentialist is just how claims like ‘water= H<sub>2</sub>O’ are supposed to be necessary: are they necessary in the same way as 2+2=4, or are their different kinds of necessity? For discussion, see (Fine 2002), who argues that there are three fundamental, distinct kinds of necessity.

means. But here their view is on very shaky ground: Kripke and Putnam argue for their essentialist externalism by means of thought experiments, by thinking about what we would say about various cases in which, say, two internally identical individuals in different environments utter the 'same' sentence or think the 'same' thought. However, if their externalism is true, there is little reason to think thought experiments like this will reveal interesting truths about the world, including that semantic externalism is true.

Given that the semantic questions we need to settle before moving to examine the mind-independent world cannot be settled by direct appeal to the mind-independent world itself, the use of thought experiments is indispensable. As a result, one cannot attempt to support a semantic thesis by appeal to thought experiments while holding a view about meaning that implies thought experiments are unreliable as a means of establishing necessity claims. That does not show that essentialism or semantic externalism are false. But if semantic externalism and its accompanying metaphysical essentialism cannot consistently be supported by the one and only indispensable philosophical method, then so much the worse for those views.

The moral of the story is that, first, the essentialist assumption about the nature of necessity is at least questionable given the lack of grounds for believing in the existence of the sort of necessary connection that underlies the semantic externalist view, because we can conceive of changes in the course of nature, as Hume indicated; second, that even if there are 'hidden connections' between kind-properties of the sort a Humean might allow, we can comfortably infer from 'conceivably p' to 'possibly p' so long as the way we understand the inference as one that concerns 'relations of ideas' and not 'matters of

fact.’ So, even if one is a committed essentialist about kinds, including water, the conceivability of water being XYZ can still entail the possibility of water being XYZ (and therefore the non-necessity of water being H<sub>2</sub>O), but this will only, for the essentialist, reveal something interesting about our *concept* of water, and nothing about kinds in nature themselves, for which the inference from conceivability to possibility will remain questionable. However, that would still show that alleged necessary connections between kind-properties that are undiscoverable by the conceivability test are likely, in principle, simply undiscoverable: one cannot justifiably say anything about them, besides that it is possible that they exist.

I said near the beginning of this section that there is a conceivability-possibility problem for my internalist view, too, but that it’s less drastic than the problem for the externalist. The way the problem arises for an internalist is this. I deny that our concepts are always easy to inspect, and I devoted some time to saying what it often takes to discover the content of one’s concepts. It is often difficult to figure out how one thinks about things. As a result, one might conceive of a mind existing without any physical things existing that is identical to the mind, or which contains the mind, or whatever; but one might abandon the idea that minds can exist without bodies upon further reflection. So, as long as we are not omniscient with respect to our own concepts, we can’t simply infer from the fact that we imagine things can be such-and-such a way to the possibility that things could be that way.

But because we are able to become aware of the contents of our concepts, the method of philosophical analysis driven by thought experiments will be generally

reliable, even if fallible. Here, philosophers ought to exercise humility. It is possible for us to discover the correct analysis of our concepts, but we can make mistakes. We simply have to try as hard as we can.

\* \* \*

My burden up to this point has been to defend a conception of philosophical analysis that allows me to answer questions of the form, “What is the correct analysis of X?” from the armchair, without any appeal to work in psychology or natural science, just thinking about my own considered judgments about thought experiments, and attempting to systematize those judgments. As we have just seen, there have been quite a few challenges to that program in the last few decades, but none of them is serious enough to warrant much revision to the old methodology.

I turn now to one final challenge that has been posed recently by Timothy Williamson (2000), who claims knowledge is not analyzable, thus the best epistemologists can hope to do is identify interesting necessary truths about knowledge.

### 1.3.5 *Is Knowledge Analyzable?*

Williamson’s basic motivation for advancing the unanalyzability of knowledge is induction from past failure:

I have long been impressed by the failure of attempts to find a correct analysis of the notion of knowledge in terms of supposedly more basic notions, such as belief, truth, and justification. One natural explanation of the failure is that knowledge has no such analysis (2000, p. v).

It’s unclear whether Williamson intends for this to be an argument that knowledge is unanalyzable or if it is just a psychological report. In any case, it is clear that

Williamsonian epistemology is unburdened by a concern to offer non-circular necessary

and sufficient conditions for knowledge. Well, fine. But does Williamson give us any reason to think the project of analyzing knowledge is misguided or doomed to failure?

Surely an inductive argument from past failure does not entail that future attempts to analyze knowledge will fail. It may just as well be that no one has quite gotten things right, or that someone has gotten things right but others have failed to appreciate it. Indeed, it is probable that epistemologists have been on the right track for a long time: for instance, most have thought that, whatever else is necessary for knowledge, *belief* is. One couldn't know that p without believing that p. And surely one cannot know that p without p being *true*. Few epistemologists have rejected belief and truth as necessary conditions for knowledge. Any inductive argument to the unanalyzability of knowledge, then, must be using as an inductive base *complete* analyses of knowledge: belief + truth + x, for some specific x, and arguing that the *conjunction* of those factors has so far failed to provide the correct analysis of knowledge. But one might just as well, from the same inductive base, see evidence of progress. We've identified *some* necessary conditions, and might as well be optimistic that we will eventually identify x. Inductive considerations do not provide compelling grounds to abandon the analytic project.

Williamson offers other considerations against the analytic project, too. He argues that knowledge may have necessary conditions including belief even if knowledge does not decompose into belief plus other ingredients that, when combined, equal knowledge (3). It's important to appreciate the exact force of his argument. The conclusion that follows is that it's difficult to argue for an analysis in the sense of a conceptual decomposition; but that leaves untouched the possibility of attempting to

identify connections between one concept of philosophical interest and others. Indeed, Williamson is rather clearly engaged in the latter sort of analysis, for he recognizes several necessary conditions on knowledge. He only denies that these necessary conditions combine to form a sufficient condition. In that way, he holds that knowledge is primitive. G.E. Moore was no opponent of philosophical analysis for holding that *good* is primitive: Williamson's claim that knowledge is not decomposable into more fundamental concepts is not a threat to the project of philosophical analysis.

The question of whether the concept of knowledge decomposes into individually necessary and jointly sufficient concepts has to be determined by doing epistemology, that is, by searching for an analysis of concepts of epistemological importance.

#### **1.4 Chapter Summary**

The argument of this chapter established a few important conclusions. These include a defense of traditional 'armchair' philosophy, a commitment to a broad kind of internalism about meaning, and that concepts are the objects of philosophical analysis. We can do philosophy by reflecting on what we're inclined to say about certain hypothetical scenarios because when we do that we are considering whether or not our concepts fit the scenario described. And doing that can yield true philosophical conclusions because there are no relevant further facts besides which concepts are the correct ones to apply such that those further facts could falsify our concept-driven judgements. As far as analysis is concerned, when one has the concepts right, one has the facts right.

In the remainder of the dissertation I will put some of these ideas to work. In the next chapter, I present and evaluate some arguments for the claim that when one knows a proposition, one must satisfy a very high epistemic standard: the standard of infallibility.

## **CHAPTER TWO: INFALLIBILITY AND CERTAINTY**

### *Chapter Preview*

In the previous chapter I defended some crucial metaphilosophical assumptions that I will rely on throughout the dissertation. I think that the philosopher's work in analyzing a concept is finished when the analysis accurately describes the concept. So, the truth of an analysis is not hostage to the requirement of making sure our ordinary concepts are regularly (or ever) used to express true claims or true thoughts. That is a good thing, for as I will argue in this chapter, our concept of knowledge requires that in order to know a proposition, a person must have justification for being certain that the proposition is true. I also argue that familiar accounts of knowledge are too weak because they allow that one could know a proposition when one lacks justification for being certain that the proposition is true: that holds even when the accounts of knowledge I have in mind are adjusted so as to make it in some sense impossible for the knower's belief to be false according to those accounts. The problem is that these other views of knowledge misidentify the relevant sense of "impossible" in "if one knows that some proposition is true, it must be impossible that the proposition is false." The relevant sense of impossibility is having justification for being certain. Philosophers have offered many different takes on what certainty is and when one is justified in having that attitude. I will argue that the relevant standard of justified certainty is the standard of infallibility. These are high standards for knowledge. In chapter three I discuss the extent of knowledge that is possible under these high standards.

## 2.1 Knowledge Requires Certainty

Epistemologists with views as opposed as Descartes and Dretske have agreed that knowledge requires certainty. The view has a distinguished history. Often, what these philosophers mean by “certain” serves to sharply separate their views. However, many have wanted to offer analyses of knowledge that allow them to say there is no knowledge without certainty, whatever that might come to on their view.<sup>31</sup>

There are lots of arguments that can be given for why we should think knowledge requires certainty. (Again, whatever “certain” means.) One might argue as Descartes did in the *Meditations* that the uneliminated possibility of imagined radical deception destroys knowledge. Or, one might point to the fact that we say you never know that a lottery ticket is a loser on the basis of the odds, apparently for the reason that it is *possible* that the ticket is, indeed, a winner (Cf. Hawthorne 2004). Or, one might argue from the apparent strangeness of sentences like, “I know that p, but I am somewhat doubtful that p” (Unger 1971, 1975).<sup>32</sup> Yet another alternative is to claim that knowledge is incompatible with beliefs that are true by luck: so to know, one must eliminate all relevant accidentality that could undermine one’s epistemic position.<sup>33</sup>

---

<sup>31</sup> I will have much to say later about the connections between one being certain that p, it being impossible that p is false, p’s being infallible, and one’s having infallible justification for believing that p.

<sup>32</sup> As I said above, the premise is that this sentence is strange: “I know that p, but I am somewhat doubtful that p.” Of course, the rest of the argument would explain why the strangeness is best accounted for by a skeptical view. The idea Unger eventually pushes is that when one knows that p one is absolutely certain that p, and when one is absolutely certain one is not at all in doubt as to whether p.

<sup>33</sup> That is one take on the basic intuition behind the famous arguments in (Gettier 1963). See (Pritchard 2005) for discussion, and for an argument that knowledge-undermining luck is a matter of there being nearby possible worlds in which one forms a false belief upon the same basis that, in the actual world, one formed one’s true belief. He argues from there to the thesis that a safety principle can eliminate knowledge-undermining luck and thus preserve knowledge. I argue in (Stoutenburg 2015a) that Pritchard’s

I think that those kinds of arguments are all good arguments for a certainty requirement on knowledge. However, I will not focus on those arguments. Instead, I will focus on an argument that attempts to establish directly that knowledge requires the elimination of every possibility of error. One reason I like this argument is that it explicitly connects epistemological theorizing to what competent speakers are likely to understand by the use of a given expression, and so my metaphilosophical work in the previous chapter is brought to bear especially easily in the context of this kind of argument. The argument is this.

- (1) Explicit statements of the thesis that one can know that a proposition is true when one cannot eliminate the possibility that it is false strike us as contradictory.
- (2) If (1), then to know that a proposition is true, one must be able to eliminate the possibility that the proposition is false.
- (3) Therefore, to know that a proposition is true, one must be able to eliminate the possibility that the proposition is false.

The argument is an instance of modus ponens, so it is valid. One might challenge the argument on the grounds that (1) is false. I will only briefly address that concern because I think (1) is rather obviously true. One might, more plausibly, challenge the argument on the grounds that (2) is false.<sup>34</sup> Here, one might argue, following Dougherty and Rysiew (2009, 2011), that while statements to the effect that one knows a proposition but

---

characterization of luck is inadequate. A result (not argued in that paper) is that safety is insufficient to eliminate knowledge-eliminating luck.

<sup>34</sup> One objection to (2) that I will not pursue in detail is this: one cannot infer from the fact that speakers take concessive knowledge attributions to be contradictory to the fact that they are contradictory, rather, speakers are just mistaken. I reject this line of argument on the basis of the view offered in the first chapter. Once we are clear on what we think knowledge is, we have discovered what knowledge is. At that level, there is no possibility that we have the correct analysis of knowledge while we are still incorrect about what knowledge is.

one cannot eliminate an error possibility indeed *sound* false, they are often true: their grating sound is to be explained by pragmatic conversational factors rather than by the hypothesis that such statements express necessarily false propositions. I will address this pragmatic explanation in more detail.

David Lewis said,

[I]t seems as if knowledge must be by definition infallible. If you claim that S knows that p, and yet you grant that S cannot eliminate a certain possibility in which not-p, it certainly seems as if you have granted that S does not after all know that p. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory. (1996, p. 549)

Lewis goes on to present a view that pays lip service to the idea that knowing that p requires that it be impossible that not-p: specifically, he devises a set of rules that say what does and does not count as possible *in a certain context*. That is why he goes on to say,

I may properly ignore some uneliminated possibilities; I may not properly ignore others. Our definition of knowledge requires a *sotto voce* proviso. S knows that p iff S's evidence eliminates every possibility in which not-p—Psst!—except for those possibilities that we are properly ignoring (554).

That Lewis found himself led to adopt this position shows how forcefully he was pulled in opposite directions by the requirement of infallibility on one side and antiskepticism on the other.

Let us consider Lewis' claim that this sounds contradictory: "S knows that p but there is some q that is incompatible with p that S cannot eliminate." Sentences with this

form—so-called concessive knowledge attributions—sound unacceptable.<sup>35</sup> I agree with Lewis that this is a datum epistemologists have to acknowledge. One cannot simply dispense with the fact that concessive knowledge attributions sound *necessarily* false. In case anyone doesn't 'hear' the grating sound, try an example sentence that fits the above structure.

1. I know that the bank will be open on my way home from work tonight, but it might close before I get there.
2. You know the bank will be open on your way home from work tonight, but it might close before you get there.
3. I know my car is parked in the Library lot, but someone might have stolen it already.
4. You know your car is parked in the Library lot, but someone might have stolen it already.
5. Steve knows he's looking at a zebra, but he can't be sure it's not a disguised mule.<sup>36</sup>

These sentences all sound contradictory, just as Lewis claimed.

---

<sup>35</sup> Concessive knowledge attributions are sentences of the form, "S knows that p but it is possible that not-p." Concessive knowledge attributions are similar to so-called abominable conjunctions. An abominable conjunction says a subject knows a proposition is true but does not know a proposition that is obviously entailed by the allegedly known proposition. Here is an abominable conjunction: "I know that I have hands but I do not know that I am not a handless brain-in-a-vat." Abominable conjunctions arise from denying the principle that knowledge is closed under known entailment, a principle I address in 3.3.5. Concessive knowledge attributions are similar to abominable conjunctions in that both are uses of language that sound grating to ordinary speakers and seem to suggest substantive epistemological doctrines.

<sup>36</sup> It would not be the first time a zoo disguised a mule to look like a zebra:  
<http://news.bbc.co.uk/2/hi/8297812.stm>.

The fact that concessive knowledge attributions *sound* contradictory does not entail that they *are* contradictory, and thus that there is no such thing as fallible knowledge. To argue for that claim, we need to consider different hypotheses about *why* these sentences sound contradictory. *One* hypothesis is that fallible knowledge is impossible: concessive knowledge attributions sound contradictory because they are contradictory. Another hypothesis is that there is something *pragmatically* defective about making concessive knowledge attributions: while concessive knowledge attributions are often true, conversational norms give rise to a kind of infelicity that makes them sound contradictory.

Trent Dougherty and Patrick Rysiew have defended a pragmatic solution to this problem (Dougherty and Rysiew 2009, 2011). They say concessive knowledge attributions sound contradictory but, in fact, are not. Their proposed solution is this. In normal circumstances, saying “It is possible that p” conveys to one’s audience that one does not know that p. That is because, in normal circumstances, if we are wondering whether or not p is true and we ask someone whether or not p is true, the subject should answer thusly: if he knows that p, he should say, “p,”; if he knows that not-p, he should say, “not-p.” As a rule of conversation, we expect speakers to ‘assert the stronger’: to answer questions with as much information as possible (without going overboard with detail) (Grice 1975). Because that maxim is true, if our speaker answers our question about p with “It is possible that p,” we will infer that he does not know that p and he does not know that not p: for, if he did know one of those things, we would be reasonable in thinking that he would have said as much. So, what is strange about “S knows that p but

it is possible for S that not-p” is that the second conjunct *pragmatically implicates* that the first conjunct is false, although the second conjunct does not *logically entail* that the first conjunct is false.

The account Dougherty and Rysiew offer here initially appears compelling. But the appearance is misleading. Dougherty and Rysiew are right about how implicatures are generated: certain uses of expressions in certain contexts are taken to convey some information beyond which those expressions literally state. For example, if I say, “I don’t like Smith,” my audience will usually at first understand me as having *said* that I *dislike* Smith. I didn’t say that: I merely claimed not to *like* Smith.<sup>37</sup> Why would anyone have ever ‘heard’ “I don’t like Smith” as expressing “I dislike Smith”? Because in normal contexts *I dislike Smith* is what speakers *intend* for their audience to hear from “I don’t like Smith,” and in general our communicative intentions are successful. That explains why “I don’t like Smith” would be heard, in a normal conversation, as “I dislike Smith.” It also explains why “It is possible that not-p” would be heard as *I don’t know that p*: because, in normal contexts, *I don’t know that p* is what a speaker would try to convey with “It is possible that p”. For, again, if I am playing by the normal rules of conversation and I know that p, I ought to ‘assert the stronger’ and say “p” when asked about p, just as Dougherty and Rysiew argue.

So, the Dougherty-Rysiew pragmatic explanation successfully explains, I think, why we would initially, uncritically hear the second conjunct in a concessive knowledge attribution as contradicting the first. But, concessive knowledge attributions *continue* to

---

<sup>37</sup> The example is from Richard Fumerton.

sound contradictory even after the pragmatic explanation is offered. Implicatures are not normally this way. It only takes a little coaxing to convince my audience that I didn't literally *say* I dislike Smith. They can easily be brought to see the difference between the semantic content of the utterance and the implicature, even if they missed it the first time around. Through this sort of clarifying procedure ("I didn't really *say* that p, I only *said* that q..." etc.) we distinguish, in actual cases, what is said from what is implicated. However, concessive knowledge attributions continue to sound contradictory. Consider how a correction speech would be received after the making of a concessive knowledge attribution: "I didn't *say* I don't know that p: I just said that I might be wrong about p. In fact, I do know that p." The grating sound only intensifies. As a result, while the Dougherty-Rysiew pragmatic explanation suffices to explain how the "I don't know that p" implicature could be generated from "It is possible that not-p," it fails as an explanation of why concessive knowledge attributions sound contradictory. That is because it cannot account for the contradictory sound of a concessive knowledge attribution made in the context of an explanation of why a previous concessive knowledge attribution was not contradictory. It cannot do all of that in a way consistent with how implicature normally works, as illustrated by the "I dislike Smith" example.

There is no standard of what sentences implicate in contexts beyond what competent speakers hear and intend to convey by their uses of language. If competent speakers overwhelmingly hear some sentence S as conveying "p" and they continue to hear "p" even after challenges are raised of the sort outlined in the last paragraph, then S *says* "p" and does not merely implicate it. One cannot defend one's preferred

philosophical positions by simply claiming that “p” is merely an implicature and not literally said by S. If competent speakers wouldn’t hear it that way, then it is not that way.

A much simpler and more compelling explanation of why we hear concessive knowledge attributions as contradictory is Lewis’ explanation: they *are* contradictory. Sentences of the form “S knows that p, but it is possible for S that not-p” are never true because our concept of knowledge holds that if S knows that p, then it is impossible for S that not-p. In that sense, knowledge requires certainty.

Well, knowledge requires *some kind of* certainty. “Certainty” has not been analyzed yet. Strictly, what the contradictory nature of concessive knowledge attributions shows is that for a person S to know that a proposition p is true, it must be impossible *for S* that p is false. But just what this comes to is a matter of debate. In the next section, I take a few steps toward identifying the relevant sense in which one must be certain that a proposition is true in order to know it.

## **2.2 In Search of Certainty**

If S knows that p, then it is not possible for S that not-p. A few crucial issues have to be addressed before we are clear on what this claim amounts to: specifically, we want a specification of the scope of “not possible” and an analysis of the type of modality indicated by “not possible”. We will see that the issue of scope is best addressed in the course of addressing the issue of modality.

### 2.2.1 *Answering Questions*

As I discussed in chapter one, a very important tool for philosophical analysis is figuring out systematic ways in which we are inclined to use language. If we discover that competent speakers are disposed to respond to certain X-related prompts in regular ways, the existence of that disposition constitutes extremely strong evidence that the speaker's standard of X is underwritten by that disposition. There is a characteristic way of answering questions about what we think we know that strongly suggests that we take ourselves to be in a position to answer such questions decisively. I will now try to highlight this phenomenon. What is significant about it is that our disposition to answer decisively questions about what we do and do not know very strongly suggests that when we know that a proposition is true, we are in a position to somehow recognize that fact.<sup>38</sup>

When we are asked questions that we do not know the answer to, we say, "I don't know." We might add some remarks about considerations we are aware of that could be relevant: not to whether or not we *know* the answer, since that question was answered with "I don't know," but to what would have been known had the answer been "Yes." For example, if someone asks me what time it is, and I don't now know what time it is, nor do I have an easy way of figuring out what time it is, I will answer, "I don't know." If there is an easily-enough accessible way of figuring out what time it is, and I know what that way is, I might add: "...but there's a clock on the wall around the corner."

---

<sup>38</sup> "Recognize that p" often conveys "know that p"; as I will argue in section 2.2.5. However, I deny that being in a position to decisively answer questions about what one allegedly knows implies knowing that one knows.

Of course, if I had known what time it is, then upon being asked, I could have said, “Yes, it’s eleven-thirty,” or just, “It’s eleven-thirty.” Interestingly, in this way we see that questions of the form, “Do you know what-X?” or “...where-X?” or “...how-X?”, while stated as questions about what a person knows, are often intended as requests for information about X: whether or not the person being asked the question *knows* the relevant X-fact is secondary. The person who answers the X-question suggests that he or she is in a position of authority with respect to the matter of X in the sense that he or she has the relevant information about X.

If I am asked what time it is and I immediately reply without so much as checking my watch, “It is eleven-thirty,” the person who asked me will likely think that I am wrongly representing myself as having the information about what time it is when in fact I do not know what time it is. I might clarify by saying, “I checked just a second ago.” That would assure my friend that I had checked the relevant source to confirm what I said was the case. Then again, I might have just given my best guess as to what time it is. That would not relieve my friend’s incredulity. Why not? Because while there are some questions that we take ourselves and others to be able to answer without further verification, perhaps involving checking some source, questions about what time it is are not that sort of question. Many questions are like questions about the time in that way.

Whether or not the person answering the question does *know* whether or not X is sometimes itself the subject of the question: “I know where X is. Do you?” In that question, the issue of X is not itself in doubt because the asker takes herself to know where X is already. What is being questioned is the *knowledge* of the person being asked

the question. When we are asked questions about what we know (rather than about whether such-and-such is the case, the distinct question I pointed to above), we are always ready to answer with a “yes” or “no.” Questions about what we know are not like questions about the time. We can answer them immediately. The fact that we take ourselves to be able to do that suggests an important feature of our concept of knowledge: that when we know that *p*, we are in a position to become aware of *that* fact.

While this might sound dangerously close to endorsing the principle that if *S* knows that *p*, then *S* knows that *S* knows that *p*—the so-called KK principle—it does not. It implies only that we are able to become aware of what we know in the sense that we are in a position to answer decisively questions about what we do and do not know. One might be tempted to describe being in a position to decisively answer questions about what we know as “knowing that we know.” I think characterizing being able to answer knowledge questions in that way is unnecessary. That is because there is an important distinction between *being aware* of one’s being in a mental state *M* in a factive way and *knowing* that one is in *M*: while knowing that I am in *M* is itself a factive relation due to the obvious fact that knowing that *p* entails that *p*, being factively aware of being in *M* does not imply knowing that I am in *M*. In light of this distinction, then absent some argument that being factively aware of being in *M* requires knowing that I am in *M*, there is no reason to think my claim that we are in a position to decisively answer questions about our knowledge implies the KK principle. (I say more about epistemic levels in section 2.2.5)

Actually, we do not always take ourselves to be able to decisively answer questions about our knowledge *at precisely the moment they are asked*: we regularly say things like, “I swear I know this!” when we find ourselves stumped as we try to recall some bit of information. Taken at face-value, “I swear I know this!” suggests ‘I know this’. But if I really take myself to know this, and the condition just stated above is correct, then I should be aware of knowing it. So, there are cases where we find ourselves inclined to say that we know some proposition, but can’t easily become aware of our knowing it.

This fact requires emphasizing the idea that when one knows a proposition, one is *in a position* to answer decisively questions about one’s knowledge. The relevant notion of decisive answerability is dispositional: when one knows that p, one is *in a position* to become aware of one’s knowing that p. It is possible to know that p while not at some time or other being aware of one’s knowing, so long as one is still disposed (at that time) to become aware that one knows. Necessary conditions invoking claims about what would be the case if such-and-such conditions were met have to say something about those conditions. My claim is that when one knows that a proposition is true, one is in a position to answer decisively questions about one’s knowledge of that proposition. Being in a position to answer such questions amounts to this: by thinking carefully, one could easily become aware both of all of one’s evidence for p and the connection between the evidence and p. There would be nothing else one would need to check, no further sources to verify. Everything that is necessary and sufficient to determine whether or not one knows that p is ‘there’ for one to find upon inspection. Whatever exactly the correct

analysis of knowledge turns out to be, it has to respect the requirement that when we have all of our evidence and its relevance to our belief in focus, we are able to answer decisively questions about what we do and do not know.

This requirement helps us restrict the relevant scope and modal status of “not possible” in the requirement on knowledge that when one knows that  $p$  it is not possible for one that  $\text{not-}p$ . The relevant scope and modal status of “not possible” must be such that when  $S$  knows that a proposition is true and  $S$  is aware of all of  $S$ 's evidence for believing the proposition and the connection between the two, it is impossible for  $S$  that the proposition is false. Knowledge eliminates error possibilities. The conditional holds in the other direction, too: when it is impossible for  $S$  that the proposition  $S$  believes is false, then  $S$  knows that the proposition is true. Eliminating all error possibilities gives knowledge. We can say now that when one knows that a proposition is true, there is no possibility that the proposition is false that one cannot eliminate at the time one knows the proposition.

For now, I will say that knowing that  $p$  requires that  $\text{not-}p$  be impossible for one in the sense that when one knows that  $p$ , one is in a position to decisively answer questions about one's knowledge. That is what it takes to be certain that  $p$ , in a sense I will further qualify in the next section.

### 2.2.2 *Certainty and Being Justified in Being Certain*

I used the notion of being in a position to decisively answer questions about what one allegedly knows as a way of getting to the heart of the idea that when one knows a proposition is true one must be certain that the proposition is true in the sense that one

can eliminate all possibility of error. The test I propose for figuring out whether an analysis of knowledge allows that one can know that a proposition is true when there is an error possibility that is (for one) uneliminated is to first make one aware both of all of one's evidence for believing the proposition and the connection between that evidence and the proposition, and to see then if one is in a position to decisively answer questions about what one allegedly knows. If so, then the analysis holds that one knows only when error is impossible. But we can imagine someone in this ideal circumstance for whom the evidence for the proposition does not eliminate all error possibilities, yet who, if asked, would say, "Yes, I'm certain the proposition is true." The person in this circumstance may be as confident as can be that the proposition is true, while the proposition is true and possibly even well-supported by the grounds for believing it. That is not enough to make error impossible.

Still, the person in this situation is certain that his belief is true, in a sense: there is no doubt this person actually entertains about the truth of the proposition. We might as well add that if asked, in a suggestive tone of voice, "Are you really *sure* the proposition is true? Might you be mistaken?" the person would continue to have no doubt about the truth of the proposition. What do we say about this? "He is sure, but he shouldn't be"? "He is not sure (because his reasons for believing are not good enough) but he is sure (because he is not at all in doubt)"?

To answer these questions we must distinguish being certain and having epistemic justification for being certain.<sup>39</sup> To be certain that a proposition is true is just to be confident enough that the proposition is true that one does not at all doubt that the proposition is true. If there is any doubt *at all* that a proposition is true, one is not certain of it. Like questions about our knowledge, we take ourselves to be able to decisively answer question about whether or not we are certain. If I want to know if someone is certain, I will just ask: “Are you certain?” or “Are you sure?” or, if I want to make sure my interlocutor is taking the question seriously, “Are you *really* certain?” or “Are you *absolutely* sure?” Provided that the interlocutor understood the question, there should be no need to hesitate before giving an answer.<sup>40</sup> If someone is sure that p at a time, there is no doubt for that person at that time *at all*. It might be that after a little more reflection the person says, “I thought about it: and yes, I’m sure” (or “...I’m positive”). In that case, the person was not sure that p but after some thought became sure that p.

---

<sup>39</sup> Throughout, when I speak of justification I mean *epistemic* justification. Epistemic justification is conceptually connected to truth (as I elaborate in the next chapter). It is distinct from the sense of ‘justification’ we use to indicate when some means is an effective way to achieve an end. Call that pragmatic justification. Perhaps the classic example used to distinguish the two is this. Suppose you’ve just been given bad news from your trusted physician: you have an illness and it is improbable that you will survive. However, you learn that people who ‘fight back’ by convincing themselves they will survive are more likely to do so. In this case, you lack epistemic justification for believing you will survive, because the physician said you probably won’t; but you have pragmatic justification for believing you will survive, because you are more likely to survive if you believe you will.

<sup>40</sup> For that reason, as I’ll suggest later, hemming and hawing before giving an answer to a question about whether or not one knows or is certain suggests to the hearer that the person does not know or is not certain. Strictly, though, the condition stated above that the interlocutor understood the question and does not *need* to hesitate is important, for one might be certain that p and yet some abnormal condition obtains that prevents the person from answering immediately even though the person was certain at the time the question was asked: a sudden strong fear of answering questions overwhelms the speaker, the speaker gets momentarily distracted, etc.

One thing that might make one become certain after entertaining some doubt is recognizing that one's grounds are good enough to justify the attitude of certainty. If I first entertain some doubt about whether or not *p*, but I reflect and realize that I have no reason to doubt *p*, I may regain that extra bit of confidence needed to make me certain. When one has evidence for belief sufficient to eliminate all error possibilities contrary to the proposition one believes, one has sufficient justification for being certain. In this case, one's certainty might not actually be justified. It could be that while one is in possession of grounds for being certain that are good enough to justify being certain, one is not certain on account of one's grounds, but for some poor reason: one took a drug that makes one certain about everything, or one has developed the bad habit of being certain of everything one's favorite blogger says and the blogger said "*p*", or whatever. In these cases, we could say one *has* justification for being certain (or: "one is justified in being certain") but one's certainty is not justified, on the grounds that arguably some proper basing condition is needed that connects having justification for being certain with one's certainty so as to make one's certainty actually justified: but that issue is not important for this chapter, so I will postpone it until the next chapter.

I call the state of maximal confidence "certainty" and the state one enjoys that justifies certainty "being justified in being certain." Some epistemologists make the same distinction, but they think of these as two kinds of certainty: these are sometimes called "psychological certainty" or "subjective certainty" on the one hand and "epistemic certainty" or "objective certainty" on the other. I think my way of drawing the distinction is preferable because when we want to know if someone is certain we want to know how

confident the person is. We will likely think the person is somehow blameworthy for being maximally confident if the person lacks justification for being certain, but that is, strictly speaking, a different issue. Furthermore, ‘certainty’ and ‘being sure’ are ordinary concepts that find regular use in ordinary language, while ‘epistemic certainty’ sure sounds like a technical expression. Finally, the two distinct notions do not seem much alike—one having to do with the connection between one’s evidence and the truth of what one believes and the other having to do with confidence—so it is hard to see how the two concepts could be two kinds of the *same* general concept ‘certainty’.

While it is possible to mount a skeptical argument that aims directly at the requirement of being certain that a proposition is true (Unger 1971, 1975)—claiming that knowledge requires certainty and that we are almost never confident enough in our beliefs to qualify as being certain of them—only the notion of having justification for being certain is necessarily connected to the truth of one’s belief and the truth of one’s belief alone. As the distinction between being certain and being justified in being certain shows, states of confidence can (sadly) vary independently of how well-supported by evidence one’s beliefs are. Insofar as we want to eliminate the possibility of error that keeps us from knowledge, what is relevant isn’t whether or not one is certain or disposed to be certain, but whether one *ought to be* certain: and that is just the issue of whether or not one has *justification* for being certain.

Although it is worthwhile to seek certainty that one’s beliefs are true, what makes it worthwhile is the thought that one’s certainty has an adequate justifying foundation. So, while I have been so far examining the sense in which knowledge requires certainty,

the distinctions I have just presented suggest that the requirement is really that when one knows that a proposition is true, one has justification that supports the attitude of being not at all in doubt that the proposition is true.<sup>41</sup> Having adequate justification for being certain that p consists in being able to eliminate every possibility that not-p.

Now to combine the conditions I have been discussing. When one knows, and one is in the ideal case of being actually aware of one's evidence and the connection between one's evidence and the proposition it supports, one is in a position to decisively answer questions about what one knows: as a result, one will be in a position to assess whether or not one has evidence that justifies being certain that the proposition one believes is true by eliminating any possibility that what one believes is false.

In the next section, I consider several influential analyses of knowledge and argue that they all fail to respect the conditions I have been arguing for. In each case, I demonstrate that the account of knowledge I am criticizing allows that one can know that p, but one would not be able to answer decisively questions about the truth of the proposition one (by hypothesis) knows. As a result, there is for one an uneliminated possibility of error.

---

<sup>41</sup> To repeat what I said above, I am not saying that one can know that p regardless of whether or not one is certain that p. I am only saying that I will focus on the requirement for knowledge concerning possibilities of error rather than the requirement for knowledge concerning how confident one must be.

Similarly, A.J. Ayer argued that knowledge requires being sure that the proposition one believes is true, that the proposition is true, and that one have a "right to be sure"—one's being sure has to be supported by a grounds that make one's being sure justified (Ayer 1956). Like Ayer, I am persuaded that the condition for knowledge that is of central importance is the one that concerns the grounds for having such a strong positive attitude toward a proposition, not the requirement of a strong attitude itself.

### 2.2.3 *A Few Defective Analyses of Knowledge*

In this section I consider several analyses of knowledge and argue that they fail because they do not satisfy the conditions on knowledge I have been advocating. The purpose of the discussion is to highlight the central role in our concept of knowledge that is played by the thought that when one knows that a proposition is true, one is able to become aware of this fact. While I will be considering paradigmatic versions of views actually defended by epistemologists, I will, where relevant, offer amendments to the proposed analyses to see how defenders of the view might try to accommodate the intuitions that lead me to reject the view in question. Of crucial importance is how defenders of the views discussed here might amend their positions to eliminate possibilities of error. When one knows that a proposition is true it is impossible for one that the proposition is false: seeing how the views I will discuss might try to accommodate this requirement will go a long way toward explaining the relevant understanding of impossibility of error.

#### 2.2.3.1 *The Gettier Problem*

Many now well-known analyses of knowledge originated as attempts to solve the so-called Gettier problem. As a result, introducing these analyses will require introducing the Gettier problem. In his famous 1963 paper, Edmund Gettier proposed two thought experiments with exactly the same structure in order to refute all analyses of knowledge that combined the conditions that a known proposition be believed, true, and well-supported (but not entailed) by the knowing subject's grounds. Call these analyses 'fallible justified true belief analyses of knowledge'. Here is one of his thought experiments (paraphrased).

Smith and Jones both applied for a job. Smith has evidence that this is true: (1) *Jones will get the job and Jones has ten coins in his pocket.* That claim entails this one: (2) *A man who will get the job has ten coins in his pocket.* Smith comes to believe (2) on the basis of recognizing that (1) entails (2). As it turns out, (2) is true, but not because (1) is true. Actually, Smith will get the job, not Jones, and Smith also has ten coins in his pocket.

This thought experiment is designed to show that fallible justified true belief analyses of knowledge are insufficient.

My interest in this section is not with solutions to the Gettier problem as such. I introduce the Gettier problem only to set the stage for the analyses of knowledge that I will consider. I am not concerned with assessing the following analyses of knowledge for success or failure at solving the Gettier problem.

#### 2.2.3.2 *No False Lemmas*

One early proposed solution is that one's basis for knowing that p must not contain any false premises. Michael Clark (1963) made that proposal shortly after Gettier presented his cases.<sup>42</sup> In the presentation of Gettier's example above, Smith's reason for believing (2) is (1). (1) is false. So, one proposed solution to the Gettier problem is requiring that propositions that serve as one's premises in an inference all be true. Where one gains knowledge by inference from premises, the premises that serve as the basis for the inference must be true.<sup>43</sup>

---

<sup>42</sup> Famously, Bertrand Russell considered what is, in effect, the Gettier problem in his much earlier *The Problems of Philosophy* (Russell 1912). There, he also offered roughly the same solution that Clark later did. However, Russell did not seem to appreciate the generality of the problem Gettier later pressed more forcefully.

<sup>43</sup> As I will discuss shortly, the no false lemmas analysis is most plausibly understood as requiring that no *essential* premises are false.

That would, indeed, block Gettier's counterexample. But let us consider an alternative scenario. This time, everything is just how it was in Gettier's scenario, except now Jones gets the job and Smith doesn't. Now, when Smith concludes that a man with ten coins in his pocket will get the job, that inference is not only believed, true, and fallibly justified, but it is an inference from all true premises. However, if Smith were asked, "Do you know that a man with ten coins in his pocket will get the job?" prior to the time at which Jones actually gets the job, he would not be in a position to decisively answer that question. He might *think* he could—he might point to Jones and say, "Yes, *he's* going to get it and *he* has ten coins in his pocket, I just saw him count his coins and I heard the hiring manager say he'll get the job"—but upon reflection he should retract this. The reason Smith is not in a position to decisively answer the question about what he allegedly knows is that there is a condition whose obtaining is part of the basis of Smith's knowledge, but Smith is not aware of that condition obtaining. The condition is that all of Smith's reasons for believing the man with ten coins in his pocket gets the job are true. That condition, though essential to the no false lemmas analysis of knowledge, is not one that Smith could be aware obtains.

Smith is not justified in being certain that the man with ten coins in his pocket will get the job. Smith's evidence may include that the hiring manager said Jones will get the job and that Smith saw Jones count out ten coins and put them in his pocket. These grounds do not eliminate the possibility that it is false that the man who gets the job has ten coins in his pocket. There are several possibilities that remain uneliminated by Smith's evidence which, were they to obtain, would make it false that the man with

ten coins in his pocket will get the job. The hiring manager could be mistaken. The background check could reveal something unsavory about Jones. Jones could fail to return his I-9 form on time. That is why, given Smith's grounds, Smith is not justified in being certain that the man who gets the job has ten coins in his pocket.

Could a strengthening of the no false lemmas view justify Smith in being certain that the man with ten coins in his pocket gets the job?<sup>44</sup> Consider: in addition to requiring the truth of the premises one explicitly thinks of when performing an inference, all of the premises of all of the arguments one would appeal to to justify those premises, and the premises for those premises, and so on, must also be true.<sup>45</sup> This stronger no false lemmas condition succeeds in eliminating a kind of possibility of error. In particular, one could not know that a proposition is true while it is possible that the proposition is false in the sense that, necessarily, when one infers a proposition from premises, it is not possible that one could come to know that the conclusion is true while even a single premise is false. So, this stronger no false lemmas condition in a way respects the idea that knowledge requires eliminating all possibilities of error.

---

<sup>44</sup> Given that the straightforward no false lemmas view is open to the objection I presented above, a weakening of no false lemmas will not solve the problem. McGrew and McGrew (2007) defend such a weakening. They say the required condition is a no *crucial* false premises condition. So, they leave open the possibility that some premise actually used in a subject's argument may be false while the subject continues to know the conclusion, so long as the premise used was not crucial for that subject to reach the conclusion. However successful the no crucial false lemmas thesis is as a solution to the Gettier problem, it should be clear that it fares no better than the straightforward no false lemmas view at granting subjects knowledge only when they are in a position to decisively answer questions about what they purportedly know.

<sup>45</sup> Clark (1963) originally defended the stronger claim that when one knows, all of the premises one used are true. Goldman (1967) criticizes it for being too strong because there can be multiple redundant grounds, and if one (or more) of those is false, it does not preclude one from having a successful chain of premises all of which are true. That is why in the sentence above I offered the idealization 'premises one would appeal to': given adequate time and reflection, one would try to make one's argument for a conclusion as strong as possible by eliminating errors in the chain of reasoning.

However, eliminating all possibilities of error *in this sense* is not sufficient for knowledge. That is because on this strengthened no false lemmas condition, knowing that a proposition is true is compatible with not being in a position to decisively answer questions about whether or not one knows. The reason the no false lemmas view is open to this objection is because the true premises that support the conclusion one by hypothesis knows need not be, according to the view, related to the alleged knower in a such a way that the knower is aware that the premises are true. The truth of the premises in the argument one uses to reach a conclusion is not a part of one's evidence, according to the view. So, it remains possible to be aware of all of one's evidence and yet not be in a position to decisively answer questions about one's alleged knowledge. As a result, the strengthened no false lemmas condition is not an improvement upon the weaker no false lemmas condition with respect to its adequacy as an analysis of knowledge.

To repeat, I am not arguing that the no false lemmas solution fails as a solution to the Gettier problem. I am arguing only that one can fail to have knowledge of a conclusion even when the premises entail the conclusion and all of the premises are true. That is because in order to know that a proposition is true one must be in a position to become aware of the satisfaction of all the conditions that constitute the truth of one's belief, and it is possible to meet the conditions for knowledge proposed by the no false lemmas view while failing my stronger condition.

### 2.2.3.3 *Defeasibility*

One might view a defeasibility analysis of knowledge as an attempt to improve upon the no false lemmas view. While the no false lemmas view says that the premises in a

justifying argument must all be true, defeasibility accounts say (roughly) that when one knows that p on the basis of some evidence, there is no true proposition that could be added to one's evidence such that that evidence plus this additional proposition fails to justify p.<sup>46</sup>

Smith has the true justified belief that the man with ten coins in his pocket will get the job. Suppose Smith's evidence for believing the man with ten coins in his pocket will get the job includes this:

(E) The hiring manager said Jones will get the job

On the basis of E plus his justified background beliefs, Smith is justified in believing:

(J) Jones will get the job and Jones has ten coins in his pocket

And on that basis, Smith is justified in believing this true proposition:

(P) The man with ten coins in his pocket will get the job.

Recall that the defeasibility condition says that when one knows that p on the basis of some grounds, there is no true proposition that could be added to one's grounds such that those grounds plus this additional proposition fail to justify p.<sup>47</sup> There is a true

---

<sup>46</sup> See (Lehrer and Paxson 1969) for defense of a defeasibility view. After offering what I have stated above as a proposed defeasibility condition, they argue that it is too strong on the basis that some evidence is misleading. I ignore that detail above because I am not focused on providing a nonskeptical Gettier-proof analysis of knowledge.

<sup>47</sup> Lehrer and Paxson reject this condition as too strong because of the possibility of misleading evidence. But if misleading evidence one does not possess can eliminate knowledge, we have very little knowledge. (Harman (1973) effectively presses this concern.) They find that result unwelcome. To see the problem, suppose that in fact the hiring manager *does* approve hires himself, so he can hire whomever he chooses. And instead of D, we could add to Smith's evidence L: someone in the lobby *said* the hiring manager cannot approve hires himself. Now, clearly E+L do not support P. So, T, while true, is misleading evidence. Lehrer and Paxson would want to say that Smith could know P if the only defeater is the misleading defeater L.

proposition that could be added to Smith's evidence such that that evidence plus this other proposition together would fail to justify P:

(D) It is not up to the hiring manager to approve hires himself.

Maybe the company that interviewed Smith and Jones requires all new hires to be approved by the Human Resources Director, not the manager of the hiring department.

While E justifies J and thereby P, E+D fail to justify P, because D defeats E. Had Smith's evidence included both E and D, he likely would have little reason (if any) to believe P. So, the defeasibility condition on knowledge says Smith's justified true belief that P is not knowledge because the true proposition D could be added to Smith's evidence, and if it were, Smith would no longer be justified in believing P on the basis of his evidence plus D. So, the defeasibility condition successfully accounts for Smith's lack of knowledge that that man with ten coins in his pocket will get the job.

It should be pretty clear that satisfying the conditions for knowledge specified by the defeasibility analysis is not enough to put one in a position to decisively answer questions about what one knows, and thus not enough to justify one in being certain that an allegedly known proposition is true. We can grant that a subject has good evidence for believing a proposition, that the subject believes the proposition on account of that good evidence, that the proposition is true, and that there is no true proposition that could be added to the subject's evidence that, combined with the subject's evidence, would fail

---

I consider the stronger defeasibility condition above for ease and because I am not concerned here with problems of misleading evidence. Since I will argue that the strong defeasibility condition is too weak of a condition on knowledge for my purposes, further consideration of Lehrer and Paxson's weaker defeasibility condition is dialectically irrelevant in this context.

to justify the proposition. Even then, the subject would not be in a position to decisively answer questions about what he allegedly knows. That is because it is possible for the subject to be aware of all the evidence he has that supports a proposition, and yet fail to be aware of something else without which he cannot justifiably be certain that the proposition is true: namely, that there are no true propositions that would defeat his justification for believing the proposition. So, for the subject, the possibility that, in fact, there is some true proposition that could defeat his justification is an uneliminated possibility of error. Uneliminated possibilities of error are incompatible with knowledge.

#### 2.2.3.4 *Causal and Nomological (Reliability) Accounts*

The next class of analyses of knowledge have a common trait: they identify a crucial condition for knowledge with a relation between a subject's belief and what is responsible for making the belief true (or highly probably true) that is determined by nature. The relation in question is a natural connection. I will consider these analyses of knowledge together. While they have differences that set them apart enough to offer different verdicts on particular cases, their similarities are more important for my purposes than these differences are.

Consider first a causal analysis of knowledge. Goldman (1967) offers an early, influential account. Like many analyses of knowledge, the causal theory is motivated in part by an attempt to solve the Gettier problem. The causal theory would deliver the intuitive verdict that in the case of Smith's belief that the man with ten coins in his pocket will get the job, Smith does not know that proposition is true because what caused him to

believe the proposition is not the fact<sup>48</sup> that makes the proposition true, but some other, irrelevant fact. Had the cause of Smith's belief been the fact that made true that *he* (not Jones) would get the job, then Smith would have known it. So, the causal theory says that when a subject knows a proposition, then in addition to the proposition being true and the subject believing it, there must be a causal connection between the subject's belief in the proposition and the fact that makes the proposition true.<sup>49</sup>

The causal analysis has a strength. Given that the subject has a belief that is causally connected in the appropriate way with the relevant fact, the belief is guaranteed to be true. In a sense, then, the causal analysis satisfies the condition that when one knows a proposition, it is not possible for one that the proposition is false: namely, that when the subject's belief satisfies the causal condition, the truth condition is thereby satisfied as well. However, that the condition is satisfied is a possibility that remains open *to the subject*. Although when the subject knows that *p* there is a fact that is appropriately connected to the subject's belief that *p* that makes it the case that *p*, the subject would count as knowing that *p* on the analysis even if the subject would easily continue to believe that *p* on exactly the same grounds even if there were no fact that made it the case that *p*. The fact that the subject gets things right in one case is often not

---

<sup>48</sup> Or whatever is appropriately causally related to beliefs.

<sup>49</sup> Goldman occasionally says the belief must be caused by "the fact that *p*," but as he is aware, this is less than precise, because it is possible to have knowledge of the future without facts about the future causing present beliefs. So, he says there must be a causal *connection* between fact and belief rather than that the fact is *the cause*, and he says the connection need not be between fact and belief, but may obtain between some common cause and both the fact and belief. He gives this example (p. 364): S can know that T will be downtown tomorrow because T said to S, "I will go downtown tomorrow." But T's being downtown tomorrow is not what provides for this knowledge. Rather, T's intention to go downtown causes both T saying to S, "I will go downtown tomorrow," and T's going downtown tomorrow.

something the subject could have figured out. As far as the subject is concerned, having knowledge will often be a sort of fluke.

Nomological reliability analyses might seem to improve on the sort of fluke that the causal analysis allows to be compatible with knowledge. Dretske (1971) and Armstrong (1973) both offer nomological reliability views. According to Armstrong, knowers are like reliable thermometers. A thermometer is reliable (in some set of circumstances) when (in those circumstances) if it is  $T^\circ$ , the thermometer reads  $T$ . A person knows that a proposition is true, says Armstrong, when the state of affairs consisting in that person's believing the proposition necessitates by a law of nature that the state of affairs consisting in that proposition's being true obtains.<sup>50</sup> Similarly, Dretske says a person knows that a proposition is true on the basis of some reason when the person would not have that reason unless the proposition were true. For both Armstrong and Dretske, the modal connection between subjects' beliefs and the truth of the propositions believed that gives rise to claims like 'the subject would not believe the proposition unless it were true' is secured by laws of nature.

The causal analysis allowed that a subject could know a proposition was true even if the subject would have the same belief under the same conditions when the proposition

---

<sup>50</sup> Armstrong is careful to insist that even reliable thermometers can sometimes get things wrong because the analogy with humans requires it. *We* sometimes get things wrong. So, if we are to often have knowledge on Armstrong's analysis, he cannot require this strong condition: 'S knows that p so long as there is a law of nature guaranteeing that S believes that p if and only if p.' That condition is too strong because it eliminates knowledge if the subject would in some circumstance believe p even if p were false. Armstrong wants to avoid skepticism. But his attempt to do so leads him to offer an account that is uninformative. Briefly, he says that a subject who is a knower must have some property such that in the circumstances a subject with that property would not believe falsely due to a law of nature. But that sounds glaringly like 'a subject knows when if the subject cannot be wrong about this sort of thing in circumstances exactly like this, the subject will not be wrong about it.'

believed was false. These nomological reliability views improve upon the causal analysis by insisting that the subject who knows a proposition must not be capable of such a mistake. They eliminate the kind of chanciness that (in part) undermines the causal analysis by appeal to necessary connections in nature.

Supposing that the strongest nomological reliability condition is satisfied, would that be enough to be in a position to decisively answer questions about what one allegedly knows? Suppose a person is a reliable believer in the right circumstances where the person's reliability will make a knowledge-relevant difference, the person believes a proposition, the proposition is true, and the person would not have believed the proposition if it had been false. Is it possible for this person that the proposition is false? It is not difficult to come up with a scenario that satisfies this description but where the person is almost completely in the dark about whether or not the proposition is true. From the subject's view, the belief is much like a lucky guess, just a lucky guess the subject always gets right in circumstances exactly like this.

Keith Lehrer's TrueTemp is an Armstrong/Dretske-knower (Lehrer 1990). TrueTemp is a perfectly reliable indicator of the temperature in the environment. Whenever he has a belief about what temperature it is in his immediate environment, it is that temperature. There are no exceptions. But from TrueTemp's perspective, there is no reason *at all* to believe the temperature is  $T^\circ$  when it is  $T^\circ$ . His temperature beliefs have the character of guesses. (And they will continue to have the character of guesses unless he comes to realize that he is perfectly reliable: but at that point, he would have evidence that entails that if he believes it is  $T^\circ$  then it is  $T^\circ$ .) He is thus not in a position to

decisively answer questions about his temperature beliefs, not justified in being certain that it is  $T^\circ$ , and unable to eliminate possibilities of error that undermine what he allegedly knows.

The nomological reliability views contain an insight that plays a central role in other analyses of knowledge. Indeed, some defenders of nomological reliability accounts point to the insight in question in places, but seem to neglect the role it could play.

Dretske (1971, p. 2) says that a reliable thermometer operating in normal conditions ‘would not have read  $T^\circ$  unless the temperature being read were  $T^\circ$ ’; similarly, Goldman’s later (1976, p. 773) reliability analysis of knowledge points to examples where a given subject lacks knowledge even when the subject’s belief is true and formed on normally good grounds because if the belief were false the subject would continue to have the belief anyway. Another set of analyses of knowledge take that insight and make it the central feature of their accounts.<sup>51</sup> That basic insight is that a subject knows that a proposition is true when the proposition is true and given the subject’s basis for believing the proposition, the subject would not hold the belief if it were false: or, if the belief were true, the subject would believe it. (Those claims are not logically equivalent, as defenders of the latter view have been eager to point out, e.g. Sosa (1999).) Perhaps the relation that determines knowledge consists in the modal connection between a subject’s

---

<sup>51</sup> Cf. Sherrilyn Roush, who writes, “Unlike, for example, the causal view of knowledge, the tracking view is indifferent to the type of process that connects a fact and the holding of a belief about that fact. As long as the proper correlational relation holds between the endpoints the process may be anything or, indeed, nothing at all. If magic does not bring knowledge, for example, on a tracking view this will be a contingent fact due to the inability of magic to underwrite the correlational properties” (2005, p. 40). She defends a type of ‘tracking view’ that she mentions in the quote.

belief, basis for belief, and the truth of the proposition believed, rather than whatever basis that modal connection has, as the causal and nomological reliability analyses wrongly insisted. I will consider such views next.

#### 2.2.3.5 *Modal Analyses: Sensitivity and Safety*

It is a bit arbitrary to separate modal analyses of knowledge from reliability analyses, both because (as we just saw) reliability analyses contained the insight that some subjunctive conditionals will be true when a belief is reliably produced, and because modal analyses require a kind of reliability in that they hold that a subject only knows a proposition is true when the subject's belief reliably tracks the truth of the proposition. Still, it is helpful to consider the types of views differently given that they analyze knowledge from different starting points: reliable belief-production grounded in nomological relations on one hand, and the truth of certain subjunctive conditionals on the other.

The most famous defender of a so-called sensitivity condition on knowledge is Robert Nozick (1981).<sup>52</sup> Like many other epistemologists, Nozick's analysis of knowledge was motivated in part by a desire to provide a solution to the Gettier problem. Nozick's diagnosis of Gettier's original case of the man who gets the job who has ten coins in his pocket locates the problem in the fact that Smith would have believed 'the

---

<sup>52</sup> Dretske (1970) also formulated what we would now call a sensitivity condition and argued that knowledge is not closed under known entailment. Like others, I focus on Nozick because his presentation of the view, though later than Dretske's, is also more developed.

man who has ten coins in his pocket will get the job' even if that proposition were false.

This led him to the following four conditions on knowledge:

- (1) p is true
- (2) S believes that p
- (3) If it were not the case that p, then it would not be the case that S believes that p
- (4) If it were the case that p, then it would be the case that S believes that p

When (3) and (4) are both satisfied, Nozick says the subject's belief "tracks the truth" of p. (3) is the sensitivity condition. After presenting his four conditions, Nozick immediately revised his view to make reference to the method by which the belief that p is formed in light of this counterexample:

A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. Clearly, we must restate our conditions to take explicit account of the ways and methods of arriving at belief. (Nozick 1981, p. 179)

The argument here is that, as stated, the grandmother's belief that her grandson is well is insensitive: if he were not well, she would believe he was well (because others would tell her that he is well). Thus, her belief fails the sensitivity condition as written. But the grandmother's belief is not obviously defective, so Nozick revised: Nozick's revision is to include in the conditions for knowledge the *method* by which one forms a belief. The grandmother's belief *based on vision* is sensitive, for if her grandson were unwell, she would not believe *on the basis of vision* that he was well. Thus the grandmother's visual belief is sensitive, although the grandmother's testimonial belief is not, because the

grandmother would believe on the basis of testimony that her grandson was well even if he were not. The revised conditions are these (p. 179).

- (1) p is true
- (2) S believes via method M that p
- (3) If p were false and S were to use M to arrive at a belief whether or not p, then S would not believe that p via M
- (4) If p were true and S were to use M to arrive at a belief whether or not p, then S would believe that p via M.

Nozick makes other revisions and qualifications, but these are not relevant to our purposes.<sup>53</sup>

Famously, Nozick recognizes that his analysis of knowledge denies the so-called closure principle, which says that knowledge is closed under known entailment.

Roughly, that principle says that if S knows that p and S knows that p entails q, then S knows that q.<sup>54</sup> It is easy to find a case of knowledge according to Nozick's conditions that fails to satisfy the closure principle. One need only find a believer who believes that p, who knows that p entails q, whose belief that p is sensitive (thus knowledge), and whose belief that q is insensitive (thus not knowledge). Here is one. I know via my visual, auditory, and tactile systems that I am in my office typing on my keyboard. If I were not in my office typing on my keyboard, then I would not believe via my senses that

---

<sup>53</sup> In particular, he is concerned with cases of overdetermination. In one sort of overdetermination case, either method is sufficient to produce the belief that p, but only one method is sensitive. The grandmother case is like that. The other sort of overdetermination case is one where the subject believes via a sensitive method that p, but if p were false, S would believe that p by a different, insensitive method. His qualifications to his account on the basis of these types of cases makes use of the concept of methods 'outweighing' one another. That notion is important to his account, but not to my purposes.

<sup>54</sup> This formulation is overly strong and thereby false, but it captures the basic idea. Hawthorne (2004, ch. 1) contains helpful technical discussion.

I am—rather, I would believe that I am wherever I would be (outside, perhaps). So, per Nozick's conditions, I know the following:

P: I am in my office typing on my keyboard.

I also know that p entails q:

Q: I am not being deceived into falsely believing that I am in my office typing on my keyboard.

However, I do not know that q, for although I now believe that I am not being deceived into falsely believing that I am in my office typing on my keyboard, if in fact I were being so deceived, I would continue to believe that I was not being so deceived. Thus, we have a case where a subject's belief that p is sensitive, the subject knows that p entails q, yet where the subject's belief that q is insensitive.

Nozick argues that the fact that his analysis of knowledge fails to respect closure is a strength rather than a weakness of his account. After all, he correctly points out, there are all sorts of things we ordinarily say we know even while we could not know various entailments that we know hold.<sup>55</sup> Although he is correct about our ordinary practice, what that really shows is that there is a problem with relying too heavily on our ordinary practice of knowledge attribution to deliver the correct analysis of knowledge. Much of what I have argued in chapter one supports that claim, and I say even more in chapter four. Here, I content myself with noting that Nozick's optimism is not widely shared: most epistemologists think that rather than inferring that the closure principle is

---

<sup>55</sup> Dretske (1970) also makes this claim.

false from the truth of his sensitivity analysis of knowledge, we should infer the falsity of his sensitivity analysis of knowledge from the truth of the closure principle.

The reason Nozick's sensitivity principle conflicts with the closure principle has to do with a feature that I haven't mentioned, namely how to analyze the subjunctive conditionals Nozick used in his account of knowledge. The sensitivity principle in particular says 'If *p* were false and *S* were to believe that *p* via *M*, then *S* would not believe that *p* via *M*'. The principle asks us to evaluate the consequent of the conditional by holding fixed a conjunction of hypothetical conditions: that *p* is false and *S* believes that *p* via *M*. Nozick doesn't say much about his preferred analysis of subjunctive conditionals besides stating that they can be interpreted using the framework of possible worlds (although he says he neither endorses nor is committed to any particular possible-worlds analysis of subjunctive conditionals (p. 174)). Following Nozick, we can interpret 'if *p* were true then *q* would be true' as true if and only if "in all those worlds in which *p* holds true that are closest to the actual world, *q* is also true... Whether or not *q* is true in *p*-worlds that are still farther away from the actual world is irrelevant to the truth of the subjunctive" (p. 173-4).

Given that analysis of subjunctives, closure-denial arises not from the sensitivity principle itself, but from evaluating the sensitivity of one proposition by reference to one set of possible worlds, and evaluating the sensitivity of a second proposition (entailed by the first) by reference to a different set of possible worlds (cf. Wright 1983a, p. 139). To see this, consider again my example of believing both that I am in my office typing on my keyboard and that I am not being deceived into falsely believing that I am in my

office typing on my keyboard. Intuitively, if I know the former, I also know the latter (by closure, because the first entails the second). But because my belief that I am in my office is sensitive and my belief that I am not falsely deceived is insensitive, by conjunction we get ‘I know that I am in my office typing on my keyboard, but I don’t know that I am not being deceived into falsely believing that I am in my office typing on my keyboard.’ The former belief is sensitive because *in worlds similar to the world in which I really am in my office*, I believe I am in my office when and only when I really am in my office. But the latter belief is insensitive, because *in worlds similar to the world in which I am falsely deceived into believing that I am in my office*, I would believe I was in my office even if I were not. The feature of Nozick’s account that gives rise to closure-failure is this shift in worlds of evaluation from worlds similar to the world in which I am in my office to worlds similar to the world in which I am falsely deceived into believing that I am in my office.

The only version of a sensitivity theory of knowledge that would even have a hope of satisfying the condition that when one knows that a proposition is true one must be able to eliminate all possibility of error would need to preserve closure, and for an obvious reason: one possibility of error that would undermine a claim to knowledge is a possibility that is entailed by what one allegedly knows. If I believe I am in my office—if I take myself to know this—then one of the possibilities I must be in a position to eliminate is the possibility that I am not in my office. For a sensitivity account of knowledge to be adequate, then, it needs to abandon the shift from evaluating antecedents and consequents by reference to different sets of possible worlds.

We could amend Nozick's theory of knowledge to make it compatible with closure by preventing shifts in worlds of evaluation between known propositions and propositions entailed by known propositions. Call a belief formed via method M 'super-sensitive' if it meets the condition that if the belief were false, under no circumstances whatsoever would the subject hold the belief via M. A tracking theory of knowledge that incorporated super-sensitivity instead of sensitivity would state the conditions in its analysis in the same way Nozick does:

- (1) p is true
- (2) S believes via method M that p
- (3) If p were false and S were to use M to arrive at a belief whether or not p, then S would not believe that p via M
- (4) If p were true and S were to use M to arrive at a belief whether or not p, then S would believe that p via M.

Unlike Nozick's account, the tracking analysis of knowledge with super-sensitivity interprets the counterfactual in (3) differently. Instead of looking only at *worlds similar to the actual world* where S falsely believes p via M, super-sensitivity says we must consider *all* worlds where S falsely believes p via M. If any of those worlds is a world where S continues to believe p via M, then S's belief that p is not super-sensitive. So, my belief formed on the basis of perceptual experience that I am in my office is super-sensitive only if it is *impossible* that I would believe I am in my office if I were not in my office. If I were a victim of radical deception—an envatted brain, or whatever—I would continue to hold the belief. So, my belief that I am in my office typing on my keyboard is not super-sensitive, and closure is preserved: it is true that if I know that I am in my office typing on my keyboard, then I know that I am not falsely deceived into believing this. Nozick himself would surely think interpreting (3) as super-sensitivity is too

strong. He explicitly claims that his subjunctive conditionals ‘if p were true then q would be true’ do not say ‘p entails q’ or that it is “logically impossible” that ‘p and not-q’ (p. 173). But, as we just saw, evaluating subjunctives his way is what gives rise to denying the closure principle, and thus to eliminating the hope of an adequate analysis of knowledge.

Sherrilyn Roush (2005, 2012) has offered an improvement on Nozick’s sensitivity account that, among other amendments, attempts to make a non-skeptical tracking theory of knowledge that is compatible with closure. Her key move is to replace Nozick’s analysis with a disjunctive analysis that explicitly endorses closure.<sup>56</sup> First, she defines knowledge by tracking. S knows that p by tracking when and only when all of these conditions are met:

1. p is true
2. S believes that p.
3. The probability that the subject believes not-p given that not-p is greater than the threshold for knowledge.
4. The probability that the subject believes p given that p is greater than the threshold for knowledge, AND the probability that the subject believes that not-p given that p is less than 1 minus the threshold for knowledge (Roush 2005, p. 45).<sup>57</sup>

---

<sup>56</sup> Another important feature of her account that is, I think, irrelevant to whether or not her account is compatible with closure is that she replaces Nozick’s subjunctives with statements of conditional probability. I am unsure whether or not that is an improvement. However, when describing how Roush’s account handles particular cases, it is still overwhelmingly natural to describe what one ‘would’ believe ‘were’ it the case that one believed something else: so even if in the final analysis Roush’s conditional probabilities should appear where Nozick had subjunctive conditionals, Roush’s conditional probabilities ought to make sense of descriptions of cases in subjunctive terms.

<sup>57</sup> This is my English paraphrase of her third and fourth conditions. Her actual notation is this:

- III.  $P(-b(p)/-p) > t$ , and
- IV.  $P(b(p)/p) > t$ , and  $P(b(-p)/p) < 1-t$

Those conditions alone are insufficient to ensure that closure holds, so she must add conditions that state what it takes to know that one proposition entails another. One knows that q entails p when:

1. q entails p
2. S believes that q entails p
3. The probability that the subject does not believe that q given that the subject does not believe that p is above the threshold for knowledge.
4. The probability that the subject believes that p given that the subject believes that q is above the threshold for knowledge.

With these in hand, Roush says S knows that p if and only if:

S knows that p by tracking p                      OR  
p is true, S believes that p, S knows that q, S knows that q implies p

Roush's account respects closure. Take the propositions I used above to illustrate how Nozick's tracking theory violated closure. Q: I am in my office typing on my keyboard. P: I am not deceived into falsely believing that I am in my office typing on my keyboard. I know that q entails p: in addition to being true, my belief that q entails p also satisfies Roush's conditions for knowing that one proposition entails another. For, the probability is very high that I do not believe (or would not believe) I am in my office typing on my keyboard if I do not (or would not) believe I am not falsely deceived into believing I am in my office typing. And the probability is very high that I do believe (or would believe) I am not falsely deceived into believing that I am in my office typing given that I believe (or would believe) that I am in my office typing. So, on Roush's view, knowledge of the entailment from this particular q to p is possible.

Clearly my belief that I am not falsely deceived into believing that I am in my office typing does not track, for it is very probable that given that if I were deceived, I

would still believe that I was not. So, if it is possible to know that I am not falsely deceived into believing that I am in my office, that knowledge will be a result of knowing that I *am* in my office and my knowledge that my knowing that I am in my office entails that I am not falsely deceived into believing that I am. Roush wants to avoid a position that is ultimately skeptical, but she accepts that offering a theory of knowledge that is compatible with closure commits her to the possibility that we sometimes know the denials of skeptical hypotheses or that we have very little ordinary knowledge.

It is familiar that if we were brains in vats we would have no indication of that. Therefore, nothing we could cite as evidence would be different (from the inside) from what a brain in a vat could cite. It is obvious that this affects our ability to give justifications for our beliefs. It is not obvious that this affects our knowledge (p. 54).

While Roush's frankness with respect to how her tracking account deals with skepticism is admirable, and while her account has the virtue of not allowing putative knowers to disregard clear error possibilities that intuitively threaten one's knowledge, her final view on the skeptical problem is unsatisfactory. We surely do ordinarily take ourselves to know many things. However, when we realize that there is something else that must be true if we really do know what we think we know, and we also realize that we do not have very good reasons to think this latter thing is true, we do not conclude that, in fact, we know this latter thing while being unable to 'give a justification' for it. Rather, we conclude that because we cannot give a justification for this latter thing, we do not know former thing. Being able to 'cite evidence' that justifies one's belief is precisely crucial for knowledge. Without being in a position to do so, one simply does not have knowledge.

While the discussion of Roush's tracking theory of knowledge helped show that sensitivity theories of knowledge are not the disaster that many epistemologists have thought they must be, many philosophers still think an alternative modal condition on knowledge is preferable to sensitivity. Specifically, they think that *safety* rather than sensitivity is an adequate modal condition on knowledge (Pritchard 2005, Sosa 1999, Williamson 2000). While the sensitivity principle says that a necessary condition of S knowing that p is that if p were false S would not believe that p, safety says (roughly) that one knows that p only when given one's basis for believing that p, p could not easily be false. Williamson connects safety and related notions like danger and risk. One is safe when there are no close possibilities of danger. A belief is safe, then, when given the method by which one forms the belief, there is no close possibility in which the belief is false. The safety principle delivers mostly the same verdicts on cases, and like the sensitivity principle, connects knowing that a proposition is true with avoiding false belief in certain similar scenarios.

Some of the main alleged virtues of safety are that it respects closure and better handles skeptical worries. As we saw with Roush, a sensitivity account of knowledge can be modified so as to respect closure because sensitivity itself does not fail to respect closure, the conjunction of sensitivity plus the standard semantics for counterfactuals does. So, respecting closure is not an advantage safety has over sensitivity. However, the versions of sensitivity that are compatible with closure threaten skepticism because they require the range of possible worlds in which one avoids false belief to be so large as to make it (actually or nearly) impossible to be mistaken when forming a belief on the

same basis that one forms one's belief in the actual world. Safety does have an advantage here: the standard semantics of counterfactuals easily lends itself to a closure-compatible condition for knowledge. Given that I believe that I am in my office typing on my keyboard, there may well be no worlds highly similar to the world in which I am in my office but where, in fact, I am not in my office. So, according to safety, I can know that I am in my office typing on my keyboard. I can also know that the fact that I am in my office entails that I am not falsely deceived into believing I am not in my office. Consequently, safety allows that I can know that I am not falsely deceived into believing I am not in my office. Safety thus allows one to know the denial of error possibilities in the same way a modified sensitivity account does.

But the problem with the way both safety and sensitivity secure ordinary knowledge such as knowledge that I am in my office is by allowing that one could know a proposition is true without having evidence that guarantees its truth. Safety and sensitivity require only an appropriate modal connection between one's belief and the truth of the belief, and that connection does not depend upon one's evidence. So, on those views I could know that I am in my office even if I am oblivious to the fact that my belief that I am in my office tracks the truth in the way safety and sensitivity say is conducive to knowledge: they grant that I could know that I am in my office while my evidence for believing that I am in my office leaves open the alternative that I am falsely deceived into believing I am in my office. If I know that p, and I know that p entails q, then I know that q. Non-skeptical accounts of knowledge that make use of safety or sensitivity use modus ponens from 'I know that p and I know that p entails q' to 'I know

that  $q$ '. But there is no reason to prefer that modus ponens to a skeptical modus tollens. If anything, there is reason to prefer the skeptical modus tollens: as I noted above, when we are faced with possibilities of error we cannot eliminate, we normally retract our original claim, rather than doubling down and insisting that we know the entailed claim as well.

Williamson argues that one's knowledge (' $k$ ') is identical with one's evidence (' $e$ ') (2000, ch. 9). A defender of the ' $e=k$ ' thesis might reply to the argument I just gave by arguing that we do have evidence that we are not deceived, indeed, we have the best evidence possible for believing this, namely that our evidence entails that we are not deceived. For when we safely believe a true proposition, we know that proposition. That bit of knowledge is evidence. And when we safely believe that a bit of our knowledge entails another proposition, we know the latter proposition. Knowledge of that entailment is also evidence. So, we may at least sometimes have entailing evidence that a skeptical hypothesis is false and that we really do know what we often think we know.

Ignoring for the moment the use of the  $e=k$  thesis in the argument, the central problem with the argument is highlighted by Roush: the person who allegedly knows the ordinary proposition and the denial of its skeptical counterpart has nothing adequate to "cite as evidence" or "give as justification" for why she thinks she knows these. It is important to note that the problem is *not* that  $S$  knows that  $p$  but  $S$  is unjustified in believing that  $S$  knows that  $p$ . We are not moving up 'levels' in this way. (I return to this point in section 2.2.5.) The problem is this: if  $S$  is not justified in being certain that  $p$ , then  $S$  does not know that  $p$ .

The concern I am raising is the same one I have been opposing to various analyses of knowledge throughout the chapter. A deep-seated intuition about knowledge is that we are in a position to decisively answer questions about what we allegedly know. So when one knows, one's evidence eliminates all possibilities of error. Williamson could reply again here that safety plus  $e=k$  allows one's evidence to eliminate all possibilities of error. This is how: when one has a safe belief, one has knowledge. Knowledge is evidence. So, when one has a safe belief, one has evidence that guarantees that one's belief is true. So, we may often have entailing evidence for the falsity of skeptical scenarios that are incompatible with what we ordinarily take ourselves to know. One might further ask, rhetorically, "What more could one want to secure the truth of one's beliefs?"

It should be rather clear that this way of securing true belief and thus knowledge fails to satisfy the spirit, even though it is compatible with the letter, of the requirements for knowledge described earlier in the chapter. So, it is time to get clearer on those requirements in order to further the account.

As we saw, we can imagine a defender of a defeasibility theory or nomological theory 'beefing up' their requirements for knowledge to eliminate every possibility of error on some construal of 'possibility of error'. Defenders of modal requirements on knowledge could do the same. Dretske, Nozick, or Roush could require that in order to know that  $p$  it must be impossible that  $p$  is false while one believes that  $p$ ; a Sosa, Pritchard, or Williamson could require that in order to know that  $p$  then, when one believes that  $p$  there is no possibility at all that  $p$  is false. The upshot of the discussion of

these theories of knowledge is that while knowledge requires the elimination of all possibilities of error, not just any *kind* of error possibility will do. For even if a modal condition were met that ensured that a particular belief would be true whenever it was formed on some basis, that would not be enough to put one in a position to decisively answer questions about what one putatively knows—unless, that is, the basis in question is the evidence the subject has.

That still leaves open a few important questions. In particular: First, given that some of the views rejected here *were* able to guarantee that a belief would be true when it was formed on a certain basis, is there more that can be said about why those views are inadequate? Second, do the necessary conditions for knowledge highlighted in this chapter imply that knowing requires being in a position to know that one knows, contrary to the argument provided in section 2.2.1? Third, what is the crucial notion of evidence relied on here?

I will answer the first and second questions in the remainder of this chapter. I save the third question for the next chapter. The result of this chapter will be clearly articulated but still schematic requirements for knowledge.

#### 2.2.4 *Internalism and Externalism about Knowledge*

In the previous section I said certain ways of securing a guarantee that one's belief is true given a certain basis are inadequate: certain guarantees are the wrong sort of guarantee. In this section, I will argue that the relevant guarantee is one that meets internalist standards for knowledge rather than (mere) externalist standards.

There are a number of ways one could approach the now famous internalist/externalist controversy in epistemology. I will not try to offer anything definitive about how to characterize ‘the’ controversy: instead I will point to a core idea associated with internalism and say that knowledge has to meet that sort of condition. I will start by attempting to say more about why internalist standards for knowledge are relevant to this epistemic inquiry while externalist standards are not. Earlier I argued that competent speakers are usually quick to give answers to questions about what they do or do not know: the presumption is that speakers take themselves to be in a position to answer such questions ‘on the spot’, without needing additional information.<sup>58</sup> It is as though we implicitly take ourselves to be in possession of everything we could possibly need to determine whether or not we know that a proposition is true. That, combined with the argument in chapter one that what competent speakers would think upon some reflection provides extremely powerful evidence of what their concepts require, implies that our concept of knowledge demands that when a person knows a proposition is true that person is in possession of everything that could be needed to determine that the proposition is known. The sense of ‘internal’ I want to explicate will capture the notion that knowers are ‘in possession’ of everything that could be needed to determine whether or not a belief amounts to knowledge.

---

<sup>58</sup> Again, the data suggest that speakers usually take themselves to be able to do this. But when considering the issue, I granted that we sometimes know a proposition is true even when we cannot recall right away everything that the proposition has going for it: the notion of decisive answerability relevant to knowledge is that with some careful reflection one could easily become aware both of all of one’s evidence for p and the connection between the evidence and p.

The conditions for knowledge surveyed in the previous section are all external in nature. Those conditions were that known propositions are not inferred from crucial false premises, that known propositions are such that no evidence could be added to the evidence one possesses in a way that would fail to justify the belief, that beliefs amounting to knowledge are causally (or otherwise nomically) connected with the facts that make the belief true, and that there is a counterfactual connection involving the holding of the belief on some basis with the belief being true. Just listing the conditions suggests a rough characterization of the common externalist feature: each of the conditions is such that one could meet the condition without being in a position to be aware that the condition is met. The suggested rough characterization is that conditions that are constitutive of knowledge must be ones that are internal *to the perspective of the knower*.

That the characterization is only a rough characterization quickly becomes evident when one considers the various ways philosophers have thought about the first-person subject's perspective. Michael Bergmann (2006, p. 13) argues that internalism must respect his so-called subject's perspective objection, which in effect requires that whatever is within a subject's perspective is something that the subject must (justifiably) conceive of as relevant to the truth of the proposition it supports. According to Bergmann, any version of internalism that requires that all justifiers be within the subject's perspective leads to vicious regress, and so (he further argues) requires abandoning internalism. So, Bergmann identifies being within a subject's perspective in

*epistemic* terms: what's in one's perspective is something one does or could have justification for believing is in one's perspective.

Other philosophers do not explicitly use the phrase "subject's perspective" to characterize their internalism, but related expressions are easy to find. Evan Fales (1996, p. 14) describes an internalist conception of knowledge as one that requires that the subject's basis for belief is within the subject's "subjective conception of the grounds" for the proposition. Conee and Feldman (2001) say internalist justification supervenes on subjects' "occurrent and dispositional mental states, events, and conditions". Earl Conee (2007) goes on to include among mental states intentional states whose existence requires the obtaining of certain contingent states of affairs in an *external* world (for lack of a better phrase).

Given how different the three positions just described are from one another, all while sounding like reasonable glosses of 'within the subject's perspective,' something more precise than that is needed to characterize the relevant sense in which I am requiring that conditions that are constitutive of knowledge must be internal. That more precise sense is this: in order to know that a proposition is true, the conditions that are necessary and sufficient for knowledge must all be facts about that with which the subject is directly acquainted. This characterization is used to avoid implying that the sort of awareness relevant to knowledge consists in anything like a first-personal judgment *that* the subject is acquainted with the conditions that constitute the truth of the proposition believed.

I briefly introduced the notion of direct acquaintance in the previous chapter. Here I will say more about it. None of what I say about acquaintance is novel, but it is necessary to make clear the internalist conception of knowledge I endorse. Direct acquaintance was first characterized explicitly by Bertrand Russell (1910) who found a central role for it in his epistemology, logic, and philosophy of mind. However, it's plausible that philosophers in the Modern period relied on the notion without explicitly characterizing it. More recently, traditional epistemic internalists—"traditional" in the sense that they accept the argumentative burdens and goals of Modern epistemology—have been more explicit about characterizing the concept and making explicit its role in epistemology: for instance, (BonJour and Sosa 2003, Fales 1996, Fumerton 1995, McGrew and McGrew 2007) all defend traditional internalist views that rely upon the notion of direct acquaintance.<sup>59</sup> Richard Fumerton has been especially committed to building his philosophical views upon the concept of direct acquaintance (1985, 1995). Perhaps the most helpful way to illustrate direct acquaintance is by example. Plausibly, there are properties of my conscious experience that I am not aware of. For example, there are all sorts of complex properties present in my visual experience, but I am unaware of exactly what they are. As I type these words, some number of alphanumeric characters are present in my visual experience, but I haven't a clue exactly how many

---

<sup>59</sup> The BonJour and Sosa reference is meant to pick out BonJour's contribution wherein he defends a traditional internalist view.

I am tempted to include Richard Feldman for some of his contributions in (Conee and Feldman 2004), especially "The Justification of Introspective Beliefs." Except for that paper, however, I do not know of any place where he describes his view as a "traditional internalist" view or anything clearly similar.

there are. The problem is not that I cannot entertain a proposition about how many there are: e.g. that there are exactly 2,205 characters, or whatever. I can do that.<sup>60</sup> The problem is that I am not able to be *distinctly aware* of each of the characters in experience at once. So, we distinguish properties of one's experience and properties of one's experience of which one is distinctly aware. This distinct awareness of a property in experience is called acquaintance. So, there are some features of experience with which we are acquainted, and some with which we are not. In this example, I may be acquainted with a few of the characters in my experience, and perhaps with the determinable property 'lots of characters', but not with all of the distinct characters at the same time.

Acquaintance (I often drop "direct") is a relation whose obtaining requires the existence of its relata. If I am acquainted with F, then F exists. If F did not exist, I would not be acquainted with F. Acquaintance is *not* usefully characterized as an *epistemic* relation, except in the sense that acquaintance is (on my view and on the view of other traditional internalists) relevant to epistemology. It is not an epistemic relation for two reasons: first, because being acquainted with some feature is not sufficient for being justified in believing or for knowing that the feature is present, and second because being

---

<sup>60</sup> An alternative view is that because I cannot pick out each of the characters at the same time, I do not have the requisite sort of concept needed to think the thought *there are 2,205 characters* rather than entertaining something that approximates that thought, such as the sentence "there are 2,205 characters". On that view, it could be granted that I am acquainted with each of the characters, but I am not acquainted with the correspondence between the characters and the robust thought in virtue of my inability to think it. Yet another alternative is that I am not acquainted either with the characters or the thought. The point is that there are a few ways of resolving this 'problem of the speckled hen' each of which involves the notion of direct acquaintance. My example is intended only to direct focus on the concept of acquaintance, not to defend my favored solution to the problem of the speckled hen.

acquainted with some feature does not consist in being justified in believing or in knowing that the feature is present. If I am acquainted with F, I can truly be said to be aware *of* F: it does not follow that I am aware *that* F is present, where this latter locution implies a judgmental sort of awareness that the acquaintance theorist rejects as a characterization of the view. Acquaintance can usefully be compared with careful attention. I am not acquainted with all of the distinct alphanumeric characters in my experience in the sense that I am not paying careful attention to each of the characters. It is also natural, but a little more misleading, to compare acquaintance to awareness: I am not aware of all the characters in my experience in the sense that I am not acquainted with them. However, I could ‘become aware’ in one sense of the phrase by highlighting all the characters on the screen and clicking the Word Count button. Then I would have a *belief* (that most epistemologists would consider justified) about how many characters there are. If someone asked, “Are you aware of how many characters there are?” I could reply, “There are  $n$  characters.” But I still would not be *acquainted* with all of the characters.

Acquaintance is the relation one often bears to one’s pain. Sharp pain has a way of grabbing one’s attention that is different from many properties of experience whose phenomenal ‘pull’ is less demanding. When I am acquainted with a sharp pain, that pain is directly before my mind. If I am directly acquainted with my pain, I am in pain. The pain is *there*, so to speak. It is extremely tempting to say that my acquaintance with pain makes it impossible for me to doubt that I am in pain. While that is often true, acquaintance with pain is only part of the story that would make my belief that I am in

pain guaranteed to be true. As I suggested in the previous paragraph, acquaintance is not an epistemic relation and it does not constitutively involve any kind of judgment, including judgments about the character of my private experience. Conceivably, then, I could be acquainted with pain while failing to believe that I am in pain: so, it is possible to be acquainted with pain while failing to believe any proposition whatsoever.

Following Fumerton, for my thought that I am in pain to be epistemically guaranteed to be true I must be acquainted with the pain, with the thought that I am in pain, and with the correspondence between the pain and the thought (1995).<sup>61</sup> When all of these conditions are met, the truth of one's thought is before one's mind. While Fumerton casts his view as a claim about noninferential justification—the kind of rational support for a belief that does not consist in having any other beliefs—the conditions he offers satisfies the conditions for knowledge laid out in this chapter. When one is acquainted with *p*, the thought that *p*, and the correspondence between the thought and fact, it is impossible for one that *p* is false, one is in a position to decisively answer the question, “Do you know whether *p*?”, and one is justified in being certain that *p*.

The central reason the satisfaction of these conditions offers the right sort of guarantee that yields knowledge is that the factors that constitute knowledge consist in relations of direct acquaintance that the proposition one believes is true. Although most of the other views of knowledge required that if *p* is known, then one's belief that *p* is

---

<sup>61</sup> While Fumerton thinks the justification one gets for one's belief when one is acquainted with the belief, the truthmaker for the belief, and the correspondence between them often results in justification that amounts to a guarantee that one's belief is true, he also thinks his view is compatible with the existence of beliefs justified in this way that are rendered very highly probable, but less than guaranteed (1995, p. 77).

related to the fact that p in some way that ‘guarantees’ that p is true, the connections in question are all ones with which the alleged knower could not be acquainted. The problem is not that an alleged knower often would be unaware of the conditions’ obtaining—after all, I do not want to rule out the possibility of dispositional knowledge—the problem is that the nature of the conditions themselves is such that neither do the conditions consist in facts about direct acquaintance, nor could one ever become acquainted with the fact that those proposed conditions obtain.

The view of knowledge suggested here is *internalist* in the sense that all factors that are necessary for knowledge are (in the ideal case) ones about relations of direct acquaintance the subject actually has to the factors that constitute the truth of the subject’s belief, or (in the dispositional case) relations of direct acquaintance the subject could easily have upon sufficient reflection.<sup>62</sup> I understand an internal condition on knowledge to be one that consists in a relation of direct acquaintance that the subject bears to some factor that constitutes the truth of the proposition the subject believes. An external condition on knowledge is a condition on knowledge that is not an internal condition. So, an internalist view of knowledge is one that holds that all conditions that constitute knowledge are internal conditions. The central problem with the other analyses of knowledge considered and rejected in this chapter is their externalist character.

---

<sup>62</sup> Like many epistemologists, I am unsure of exactly how to characterize the relevant subjunctive conditionals that would adequately analyze ‘*would become directly acquainted upon sufficient reflection*’. Also like many epistemologists, I want to offer an analysis of knowledge that does not *analytically* rule out the possibility of dispositional knowledge, even if dispositional knowledge turns out to be impossible for other reasons.

Consider the no crucial false lemmas view, for example. On that view, one can know that *p* when one believes that *p*, has strong (but not entailing) evidence for *p* in the form of some other justified beliefs that support *p*, *p* is true, and those justified beliefs in which the justification of *p* for one consists are all true. One cannot be acquainted with the last condition obtaining (at least not in the sorts of cases the no false lemmas condition is designed to accommodate, like the original Gettier scenarios). For that reason, the no crucial false lemmas condition is an *external* condition on knowledge. It is external in just the sense that one cannot be acquainted with its obtaining.<sup>63</sup>

Similarly, defeasibility is an external condition. Whatever one's evidence for *p*, the factor that gets one from a highly justified true belief that *p* is the fact that there is no evidence that would, if added to one's evidence, undermine one's justification for *p*. But one could surely never be acquainted with such a condition's obtaining.

On both of these views of knowledge, as well as the others rejected earlier, the alleged knower is cognitively isolated from truth. Early in the first chapter I described the motivation of the project of searching for an analysis of propositional knowledge as one of seeking to secure the truth of our beliefs. That basic motivation strongly supports the view that knowledge is internal. When one knows that a proposition is true, one is

---

<sup>63</sup> That is not to say the no false lemmas view is false. It may well be that an adequate conception of knowledge will render no false lemmas a necessary but redundant condition on inferential knowledge. Consider: if one knows that *p* on the basis of inference from *q*, and one is directly acquainted with the thought that *q*, the fact that *q*, and the correspondence between *q* and the fact that *q*; and *q* entails *p*, and one is acquainted with the fact that *q* entails *p*, the thought that *q*, and the correspondence between the thought that *q* and the entailment from *q* to *p*, then the no false lemmas condition would be met: there would be no crucial falsehoods in the inference from *q* to *p*. However, the reason there are no crucial falsehoods in the inference from *q* to *p* is that the truth of *p* is guaranteed by direct acquaintance with truthmakers for every proposition in the inferential chain. It is unlikely that defenders of a no crucial false lemmas condition would want to endorse such strict requirements for inferential knowledge.

directly confronted with its truth.<sup>64</sup> It is in this sense that a belief that is known to be true is one in which one is justified in being certain: when one knows a proposition is true, one is directly acquainted with the truth-maker for the proposition and the correspondence between the two. It is for that reason that one ought to entertain no doubts as to whether or not p. It is also for that reason that when one knows that p, one is in a position to decisively answer questions about whether or not one knows that p. One needs only to check one's experience and see whether or not one is acquainted with everything that constitutes p's truth.

### 2.2.5 *Epistemic Questions and Epistemic Levels*

At the end of section 2.2.3 I said I would answer two questions in the remaining part of this chapter. I just finished answering the first question: "What is the relevant sense in which an adequate analysis of knowledge must hold that someone who knows that p has a guarantee that p which eliminates all possibilities that not-p?" The second question was, "Do the requirements for knowledge described in this chapter imply that a necessary condition on knowing that p is being in a position to know that one knows that p?" I will answer that second question in this section. I will defend a negative answer.

Part of the motivation for the strong requirements on knowledge sketched in this chapter is the intuition that when one knows that a proposition is true, one is in a position to decisively answer questions about one's knowledge of that proposition. The relevant

---

<sup>64</sup> Butchvarov (1970, p. 49) expresses a similar sentiment: "[W]e continue to feel that one who knows that p does not merely have evidence, however good, that p, or reasons for believing that p, but somehow *has the truth of p itself*" (italics mine).

sort of question is, On what basis do you know that p? The answer may well be negative. When we ask a person, “How do you know?” and the person is not able to come up with a good answer rather quickly, we take this as reason to think the person does not know the proposition after all. (Now, as far as our ordinary discourse and thought is concerned the person *may* still know it, but hemming and hawing provides good reason to think the person does not know.) But assuming the answer to the knowledge question is positive, we expect to hear a reason that demonstrates to us that the person knows the proposition in question. The important feature of our knowledge-talk that I am trying to highlight is not that a person who knows can demonstrate that knowledge to others, or even to him- or herself. Rather, it is the thought that as knowers, what serves as the basis for knowledge is factors with which we are acquainted.

Roderick Chisholm (1977, p. 18) calls the following questions about knowledge “Socratic” questions: “‘What *justification* do I have for believing this?’ or ‘What justification do I have for thinking I know that this is something that is true?’” He seems to presume that the presupposition of each question is true—for the first, that there *is* a justification for believing this, and for the second, that one *does* know that this something is true—and it further seems that the questions’ having answers is what makes the questions ‘Socratic’.<sup>65</sup> He further seems to presume that these questions are asking for at least roughly the same sort of answer. However, the questions do not—at least on their surface—ask about the epistemic status of propositions on the same *epistemic level*.

---

<sup>65</sup> At least that is how I read footnote 5 on page 18.

William Alston (1989) accuses Chisholm of “level confusion” for seeming to run the question of *first*-level justification with the question of *second*-level justification.<sup>66</sup>

A proposition on the first epistemic level is one that lacks an epistemic operator. I might be justified in believing that p. Here, ‘p’ is on the *first* level. I might also be justified in believing that I am justified in believing that p. Here, ‘I am justified in believing that p’ is on the *second* level. That is because the proposition I am justified in believing has an epistemic operator as part of its content. We can think of propositions on different epistemic levels this way:

1. p
2. Jp
3. JJp

...and so on.

I say that knowing that p requires that one is in a position to decisively answer the question, “How do you know that p?” It’s tempting to identify that apparently first-level question with the second-level question, “Do you know that you know that p?”

But I do not think that being in a position to decisively answer the question “How do you know that p?” entails that when one knows that p one *knows* that one knows that p. The key notion of decisive answerability concerns the epistemic status of one’s belief that p. It does not immediately concern the conditions under which one is epistemically

---

<sup>66</sup> It is not clear that Chisholm is guilty of anything that deserves to be called level *confusion*, as Alston seems to acknowledge (p. 157). Alston further states (on page 55) but does not argue that Butchvarov (1970) makes “full-blown” level confusions. I find no support for this claim in Butchvarov’s discussion in the section ‘Knowing That One Knows’, which successfully avoids confusing levels.

warranted in answering a question about the epistemic status of one's belief that *p*. These are importantly different. If the notion of decisive answerability did concern the latter idea—if the idea were that one knows that *p* only when one can *answer* the question in the form of providing a warranted judgment to oneself or to an audience—then further questions of justification would arise regarding the basis upon which one is entitled to *assert* (or judge) that one knows that *p*. It is overwhelmingly plausible that there are some epistemic conditions on assertion, whether the condition is knowledge (as Timothy Williamson argues (2000)) or some weaker justification condition. If we think of the key notion of decisive answerability as a kind of saying to oneself or others, “I know that *p*”, *then* one is open to another question: “On what basis do you say you know that *p*?” If the first question required an answer in the form of an epistemically warranted assertion or judgment, the second does too, and you're off on a (surely vicious) regress.<sup>67</sup>

However, the reason for rejecting the construal of decisive answerability as one of actually answering questions is not the vicious regress that such a construal invites.<sup>68</sup> It is because that way of construing decisive answerability that is relevant to knowledge is simply the wrong idea. Whether or not one can demonstrate to oneself or others *that* one

---

<sup>67</sup> That being said, it may often be the case that when one is in a position to decisively answer the knowledge question, one knows that one knows that *p*. However, being in that position does not entail knowing that one knows. Furthermore, one might want a stronger notion of decisive answerability than the one I have argued for here. In particular, one might want to add that when one is in a position to decisively answer the knowledge question, one (actually or possibly) represents to oneself that one knows. Doing that would require possessing the concept of knowledge. Epistemologists concerned that an adequate account of knowledge should grant that knowledge is possible for subjects other than normal adult humans may object here that lots of knowers lack the concept of knowledge. I disagree both that an adequate account of knowledge must allow that knowledge is possible for anyone and that cognitively unsophisticated subjects obviously lack the concept of knowledge. Cf. (Stoutenburg 2015b, c).

<sup>68</sup> I have argued elsewhere that views that imply vicious epistemic regresses are not thereby false (Stoutenburg 2015c).

knows is beside the point when it comes to one's knowledge.<sup>69</sup> My preferred way of thinking about it is that when one knows that p one is *in a position* to decisively answer the question, "Do you know whether or not p?" Being in a position to do this requires nothing more than being acquainted with the factors that constitute the truth of p. That is all it takes to have the answer. When one has everything necessary to justify being certain about p before one's mind, one knows that p; that is all it takes to be in a position to answer the knowledge-question affirmatively. It is possible to be in the position of being acquainted with all of the factors that constitute the truth of p without being in a position to be acquainted with the factors that constitute one's knowledge that p. And when one is acquainted with all of the factors that constitute the truth of a proposition, there is no additional justification that would put one in a better position to be justified in being certain that the proposition is true, or that would eliminate even more possibilities of error. So, when one thinks, "Do I know that p?" the way one verifies for oneself that one does in fact know that p is not by coming to know that one knows that p, but by getting oneself acquainted with those factors in which the truth of p consists.

As a result, it is possible to know that p without knowing that one knows that p. However, it may well be that *often* when one knows that p one knows that one knows that p. To know that p requires being acquainted with those factors that constitute the truth of p. To know that one knows that p, would require being acquainted with the factors that constitute knowing that p: so, it would require being acquainted with one's being

---

<sup>69</sup> So, I am with those philosophers who, like Alston (1989, p. 166) distinguish one's *being* justified/having knowledge and one's *showing* that one is justified/has knowledge, and do not require that one is able to show that one knows in order to know.

acquainted with the factors that constitute the truth of p.<sup>70</sup> It is an advantage of a view that claims strict conditions for knowledge that if one were to have it, it may not be too difficult to ‘level-up’. Having said that, I repeat: knowing that p does not entail knowing that one knows.

### 2.3 Chapter Summary and Conclusion

I will now conclude the chapter with some final remarks and indicate where I will go with these ideas in the next chapter.

One may call “infallibilist” any view of knowledge that implies that when one meets the epistemic conditions required for knowledge one thereby meets the truth condition: in other words, the truth condition is analytically redundant because once the epistemic condition is stated in the analysis there is no reason to state separately that the truth condition is met. My view is infallibilist in that sense because I require that when one knows that p one is acquainted with the fact that p. Acquaintance is a factive relation, so if I am acquainted with p, then p. My view is also internalist in the sense characterized above: it is internalist in the sense that all the conditions for knowledge I specified are internal conditions. They are internal conditions because they consist in facts about the relations of direct acquaintance.

I will call the view of knowledge defended in this chapter and for the rest of the dissertation “internalist infallibilism about knowledge”, or “internalist infallibilism” for short. Here is a statement of the view.

---

<sup>70</sup> (Fumerton 2006a) makes a similar argument in the course of arguing that his version of internalism has a better solution to skepticism than opposing externalist views do.

S knows that p =df. (1) S is acquainted with the belief that p  
(2) S is acquainted with the fact p  
(3) S is acquainted with the correspondence  
between the belief that p and the fact p.

A few remarks are in order. Addressing them forms the content of the next chapter.

First, this is a statement of the view as it concerns foundational knowledge.

Epistemologists overwhelmingly think that much of what we know is based upon inference from other things we know. In the definition above, the stated conditions seem to exclude the possibility of inferential knowledge. I will have more to say about this in the next chapter.

Second, while many epistemologists are internalists about epistemic justification few are internalists about knowledge. The best reason (as far as I can tell) to be an internalist about justification but not about knowledge is the well-supported suspicion that internalism about knowledge implies that we have very little knowledge. As I argue in more detail in chapter four, that concern does not undermine internalism about knowledge. We find that speakers are disposed to give up knowledge claims when error possibilities are pointed out that the subject cannot eliminate. That is one of the central intuitions driving this chapter as well. The intuition points to internalism about knowledge. And as I argued in chapter one, the philosopher offering an analysis has no job besides describing the concept being analyzed. There is no further duty to make ordinary beliefs about what people know come out true. There is a challenge for my view though, and that is to offer an analysis of epistemic justification without identifying justification with knowledge. For even if there is little knowledge to be had, there is still hopefully plenty of justification to be had. At minimum, I hope I have been offering

good arguments that justify the conclusions I have defended. I will address the challenge of distinguishing justification and knowledge in the next chapter.

Third, I have said little about the notion of evidence that came up earlier in the chapter, where I said knowing that  $p$  requires that one has evidence that guarantees that  $p$  is true. In the next chapter, I will elaborate on my conception of evidence and show how having evidence is relevant both to knowledge and to justified belief.

Fourth and finally, I will describe how the view of knowledge described here should affect what we think of as the extent of our knowledge.

## CHAPTER THREE:

### EVIDENCE, BASING, AND THE STRUCTURE OF KNOWLEDGE

#### *Chapter Preview*

The focus in this chapter is on elaborating the theory of knowledge advanced in the last chapter by filling in some important details and defending some of its implications. Here, I will argue for my preferred analyses of evidence, the structure of knowledge and justification, and conclude with a discussion about the extent of our knowledge and justification. More precisely, I will argue, against the currently popular propositionalist view of evidence, that evidence for one's beliefs ultimately consists in facts with which one is acquainted, that neither knowledge nor doxastic justification requires the satisfaction of an external basing condition, that all foundational epistemic support amounts to knowledge (and not mere justification), and that our knowledge is very limited.

#### 3.1 Evidence

Epistemologists often rely on the concept of evidence without saying much about what it is. In the previous chapter, I said one knows only when one's *evidence* guarantees that the proposition one believes is true. I will explain the relevant notion of evidence more fully here.

According to the way we normally talk, on a crime scene "evidence" consists in physical objects: a bloody knife, fingerprints on a light switch, a hair fiber. For a meteorologist a reading on a barometer is "evidence" of pressure systems that bring weather changes. When asked what "evidence" we have for believing some proposition,

we will often cite as “evidence” other things we take ourselves to know or justifiably believe. These different uses of “evidence” point the way to distinct conceptions of what evidence is. The philosophical habit of seeking a unified analysis of a concept suggests that if all of these uses of “evidence” are legitimate it must be because they have some common feature. In the previous chapter I argued that what one knows, one knows on the basis of facts with which one is acquainted. More specifically, I want to defend this claim about evidence: fundamentally, evidence for one is identical to facts with which one is acquainted. That account excludes taking literally talk about physical objects, readings on barometers, and even other bits of knowledge as fundamental evidence.

The famous epistemic regress argument points the way to understanding what fundamental evidence is. One considers a belief (p) that one takes to be well-supported and asks oneself what p has going for it. One will likely think of another belief (q) one also takes to be well-supported, and which one thinks provides support for p. The very same question that arose concerning the credentials for p arises for q, as well: for if q is to justify one’s belief that p, then q must itself be justified. So we ask what q has going for it. Suppose we say what justifies q is another belief: r. Clearly, if r is to justify q which justifies p, then r must be justified, too. There are a few structures of justification compatible with this way of thinking about the dependence of justified beliefs on sources of justification. The *infinetist* says we could go on forever, never ending the regress of justifiers.<sup>71</sup> The *coherentist* says the structure could consist in links that hold between the

---

<sup>71</sup> While most philosophers would consider the infinite regress option a capitulation to skepticism, Peter Klein (1998) argues that it is the correct analysis of the structure of justification and that it does not necessarily lead to skepticism.

propositions believed such that, while the justification of every proposition consists partly in citing itself as a reason, the non-linear connections between the propositions makes the appearance of vicious circularity illusory.<sup>72</sup> Finally, the *foundationalist* says the regress of justifiers for a given person's belief must eventually terminate in propositions that are themselves justified for that person, but their justification is not constituted even in part by other propositions that are justified for that person.

The acquaintance theory of knowledge I defend is a version of foundationalism. I will not argue for foundationalism here. Rather, I will defend the conception of evidence that I think best suits an acquaintance theory of foundational knowledge. For that sort of foundationalist, fundamental evidence *cannot* consist in propositions for the obvious reason that if a person's belief is to justify another proposition, the justifying belief must itself be justified. In that case, the supported proposition would not be foundational after all. The justifying proposition would need to be supported by evidence, but if being supported by evidence consists in being supported by other justified propositions, then the view is incompatible with foundationalism.

So, an acquaintance theory of knowledge needs a conception of evidence that is non-propositionalist in the sense that it denies that evidence is to be identified with propositions.<sup>73</sup> A currently popular view about evidence is that it consists precisely in propositions (Dougherty 2011, Williamson 2000). Williamson and Dougherty give the

---

<sup>72</sup> Cf. (BonJour 1985).

<sup>73</sup> Such a non-propositionalist view of evidence is, however, compatible with the view that sensory experience has propositional *content* and that this 'given' content can serve as fundamental evidence. In that case, one's evidence consists fundamentally in sensory experience, rather than propositions that are justified or known to be true.

following arguments for their claim that evidence must consist in propositions. The first argument starts with the premise that hypotheses explain evidence: the hypothesis that I am getting a cold explains the sensation in my throat. According to a propositionalist about evidence, the evidence explained is *that I have a sensation in my throat*, not the sensation in my throat. Here is a second argument propositionalists use to support their view that evidence consists in propositions: evidence has a probability. When we use Bayes' Theorem to calculate the probability that a hypothesis is true on the basis of evidence, we assign the evidence a number that represents the probability of the evidence. But, the propositionalist continues, it makes no sense to assign a probability to anything but a proposition. So, they conclude, evidence consists in propositions.

While these arguments are suggestive, they are not conclusive. The way to avoid the conclusion is to recognize that the arguments fail to appreciate the difference between fundamental and derivative evidence. The derivative evidence I have for believing some proposition *p* will often consist in a proposition *q*. It may be that *q* is a description of something that is essentially non-propositional. It is this latter non-propositional thing that I am calling fundamental evidence. It is fundamental evidence because it is the ultimate basis for believing that *p*, beyond which there is no further basis.

Here is an example. Williamson and Dougherty say that the hypothesis that I have a cold explains the sensation I have in my throat. Here, they insist that what is explained is *that I have a distinct sensation in my throat* rather than allowing that what is explained is literally a sensation in my throat. I think we can make sense of the claim that what is explained is literally a sensation in one's throat. The distinct sensation in my

throat corresponds to the proposition *that I have a distinct sensation in my throat*. The proposition *that I have a distinct sensation in my throat* fits the sensation in a way that someone who is aware of both the sensation and the proposition is able to recognize as an accurate description of the sensation, even though the sensation itself is not a proposition.<sup>74</sup> It would be a mistake to strain the metaphor too far, but the correspondence relation in question is like the way in which a verbal description of a person can adequately ‘fit’ a visual description of a person. If I describe someone as short and having fair-colored skin, then I hand you a picture of Kobe Bryant, you would think my words did not correspond to the pictured individual as I had suggested. We recognize correspondence between propositions and things other than propositions.

In response to propositionalists about evidence we can say that in each of the examples the propositionalist offers, the evidence described is *derivative* evidence: if the hypothesis *that I have a cold* explains the proposition *that I have a distinct sensation in my throat*, that is because the evidential status of *that I have a distinct sensation in my throat* depends upon my being acquainted with the connection between that proposition and the sensation itself. After all, were it not for the sensation in my throat, there would be no plausibility at all to the suggestion that the proposition *that I have a distinct sensation in my throat* is evidence for me that I have a cold. Facts about experience, not propositions, are fundamental evidence.

---

<sup>74</sup> The move that undermines propositionalism about evidence is the same move as the one that internalists appeal to as their preferred solution to ‘the myth of the given’. See Fales (1996) and Bonjour's contribution in Bonjour and Sosa (2003).

Facts about experience are evidence. So, they help justify beliefs. But more than the mere existence of evidence is needed for justification or knowledge. In addition, the subject in possession of justification or knowledge must also possess the evidence and be aware of the connection between that evidence and the proposition the subject believes. There is a sense of ‘evidence’ not terribly relevant to epistemic justification, and in that sense facts no one is aware of make propositions probable by themselves. Recall the crime scene example. On the scene, there is a bloody knife, fingerprints on a light switch, and a stray hair fiber. If these three things are all there, but the investigator has not discovered them, the probability that a murder occurred given the evidence at the scene stays exactly the same. That is because each of those pieces of evidence is a constituent of the truth-maker for the proposition *that Jones killed Smith*. No one, including the investigator, has justification for believing *that Jones killed Smith* simply on account of the existence of evidence that supports that fact. Put another way, it is not probable *for the investigator* that a murder occurred on account of the mere existence of evidence supporting that fact.

For a proposition to be probable for a person requires in addition that the person *possess* the evidence. Suppose now that the investigator is aware of the bloody knife, the fingerprints on the light switch, and the stray hair fiber: those constituents of the truth-maker for the proposition *that Jones killed Smith*.<sup>75</sup> Now the investigator’s evidence supports believing that Jones killed Smith. But the investigator might *still* lack

---

<sup>75</sup> Feldman (2004) defends a similar view: “the evidence a person has at a time...[is] the things a person is thinking of or aware of at a time.”

justification for believing that Jones killed Smith: perhaps the investigator has failed to see the connection between these pieces of evidence and the proposition. While the proposition that Jones killed Smith is probable on the investigator's evidence, it is still not probable *for the investigator* that Jones killed Smith.

To generalize, it is possible to be in possession of evidence that makes a proposition probable or guaranteed to be true, even while lacking justification for believing that proposition. That is because the factors that constitute epistemic justification are both evidential and non-evidential. Evidence is the sole evidential justifier. On my view, two crucial non-evidential justifiers are the existence of a truth-making or making-probable connection between one's evidence and the proposition that evidence supports, and one's awareness of that connection. Other epistemologists will add to the list of non-evidential justifiers facts about the reliability of a particular process, or the sensitivity of a particular belief-forming method.

Fundamental evidence consists in facts about experience, and non-evidential justifiers include relations of acquaintance with evidence and with the correspondence between thoughts and evidence. When one knows, one's evidence guarantees that one's belief is true. What that comes to upon further investigation is that when one knows, it is impossible that the facts one is acquainted with exist while the proposition one believes is false. That is only a necessary condition for knowing. To know that the proposition is true, one must also believe it and be acquainted with one's evidence, the belief, and the correspondence between them.

That leads naturally to the next issue left open by the last chapter, which is this: in order to know that a proposition is true, how must one's evidence be related to the proposition one believes, besides corresponding to it? Must one's belief also be caused by one's evidence, or by the fact that one is acquainted with the evidence and its correspondence to the proposition one believes?

### **3.2 The Structure of Knowledge and Justification**

In this section I investigate a few consequences of internalist infallibilism for the analysis of the structure of knowledge and justification. As I will now show, some of those consequences are striking, including a denial of a basing condition for both knowledge and justified belief, and a rejection of fallibly justified foundational belief.

#### **3.2.1 *Propositional and Doxastic Justification***

Before getting to more controversial matters it will be helpful if I first explain a distinction commonly used in epistemology. That is the distinction between propositional and doxastic justification. I alluded to the distinction in the previous chapter without naming it. One has propositional justification for a proposition *p* when one meets the conditions for possessing justification for *p* regardless of whether or not one believes that *p* *as a result of* one's justification for *p*; or, indeed, whether or not one believes that *p* at all. A helpful gloss for "S has propositional justification for *p*" is "S has justification for believing that *p*" or sometimes "there is justification for S to believe that *p*."

To say that someone has doxastic justification for *p* implies that person believes *p*: not only that, but that the person's belief is 'based on' the justification the person has

for believing *p*, rather than on some other justification or none at all. So, one has doxastic justification for the belief that *p* when one has propositional justification for *p* and one's belief that *p* is 'based on' that justification in a sense of 'based on' to be explained momentarily. When someone has doxastic justification for the belief that *p* we say the person "has a justified belief that *p*" or "S's belief that *p* is justified."

### 3.2.2 *The Basing Condition*

The leading intuition behind internalist infallibilism is the notion of decisive answerability articulated in the last chapter. When one knows that *p*, one is in a position to decisively answer the question, "Do you know whether *p*?" in the affirmative, and when one does not know that *p*, one is in position to decisively answer in the negative.<sup>76</sup>

However appealing that may sound, it appears to be at odds with a widely-shared intuition that when one knows that *p*, rather than merely having knowledge-strength propositional justification for *p*, one's belief that *p* must be *based upon* the justification one has for *p*. Similarly, many think that when one has a doxastically justified belief that *p* rather than a merely propositionally justified belief that *p*, one's belief must be based on the justification one has for it. The belief that *p* would not be justified if it were based on something other than its justification, or based on nothing at all, the thought goes.

Here is an example that will illustrate why many epistemologists think basing is necessary for doxastic justification and knowledge. At the time of my writing this chapter—January 2016—many people in the United States are possessed by Powerball

---

<sup>76</sup> As I argued in the previous chapter (2.2.5), being in a position to decisively answer the knowledge question is consistent with failing to know that one knows.

lottery frenzy. Just last week the Powerball grand prize was 1.4 *billion*. Powerball play goes like this.<sup>77</sup> There is a drum with sixty-nine white balls and a drum with twenty-six red balls. The winning number is determined by five numbered white balls and one numbered red ball falling into chambers that display the numbers on the balls. Prior to the draw, players purchase a card and choose five numbers for the white balls and one number for the red ball. A winning Powerball ticket can get the first five numbers (the white ball numbers) in any order, but must get the Powerball number (the red ball number) exactly right, and on the Powerball ticket space marked ‘Powerball’. The odds of winning the jackpot are 1:292,201,338.

Suppose I hear that the jackpot is all the way up to 1.4 billion dollars and I buy a ticket.<sup>78</sup> I believe my ticket will lose. Intuitively, I have more than adequate justification *for* believing I will lose because the odds are simply staggering. If asked why I believe my ticket will lose, I will cite those statistics: “Yeah, I’m pretty sure I’ll lose—the chances of winning are 1:292,201,338!” But suppose that what caused my belief that I will lose is not my recognition of the overwhelming odds that I will lose, but the fact that I read on horoscope.com that only hard work and determination, not luck, can improve the life of a Gemini this week. We can even add to the story that I am so convinced by

---

<sup>77</sup> All of this information is from the Powerball website ([www.powerball.com](http://www.powerball.com)). The website has a page that warns users about “lottery scams” that dupe people into spending money for nothing. Perhaps it is the result of an innocent omission by the page creator that the page does not list the Powerball lottery itself among the scams.

<sup>78</sup> Indeed, I did buy a ticket and I did believe it would lose. But, as they say: Who knows? The next winner could have been me.

my horoscope that I would believe my ticket was a loser even if I thought the odds of a ticket with my numbers winning was *probable*.

Many epistemologists would diagnose that story in the following way. I have *propositional* justification for believing that my ticket is a loser, but my belief that my ticket is a loser is unjustified (that is, I lack *doxastic* justification) because it has a defective basis. If, contrary to fact, I *had* based my belief on my recognition of the slim odds of winning, then I would have a justified belief. But I based my belief on a horoscope, so my belief is unjustified.

Exactly why we should think I do not have a justified belief that my Powerball ticket is a loser is a matter of controversy. Here are two of the most popular diagnoses of why my belief is unjustified even while there is justification for me to hold it. One is that the cause of my belief—trusting a horoscope for guidance—is a cause that will usually produce false beliefs. The other is that my belief about the cause of my belief is mistaken: perhaps regardless of the causal origin of one's beliefs, one must at least believe that the cause of one's belief is legitimate.

Both of these diagnoses (and there are other possibilities) suggest interpretations of the basing condition. A crude causal basing condition would say that a belief B is based on C if and only if C causes B. On that view, my belief that my ticket is a loser is not based on its justification because the cause of the belief is the horoscope, not my appreciation of the odds. A doxastic-causal basing condition would say that B1 is based on C if and only if (C causes B1 and there exists another belief B2 which says C caused B1). In addition to faulting my belief for being caused by something other than its

justification, a doxastic-causal basing condition could also fault my second-order belief for being in error regarding the cause of my first-order belief. Finally, a plain but still crude doxastic basing condition says that B1 is based on C if and only if there exists a B2 that says B1 is based on C. According to that view, one's basis for belief is whatever one believes it is. Presumably a defender of that view will add that a belief is *properly* based on C only when the belief is based on C and C is a good reason to hold the belief.<sup>79</sup>

Even that brief sketch of possible basing conditions should give epistemic internalists pause before incorporating basing conditions into their analyses, for whether or not such conditions are satisfied could easily turn out to be the sort of thing regarding which the subject could not become aware. The causal basing condition is the easiest example. While the issue is controversial, I think we are not directly acquainted with causal relations. So, if I incorporate a basing condition on knowledge that says one knows that p only when one is acquainted with all of the factors that constitute one's knowledge that p, and one of those factors is that my belief that p is caused by the evidence I have for p, I will have thereby denied that when one knows that p one is in a position to become aware of everything that constitutes my knowing that p. That sort of concern might steer an internalist in the direction of a doxastic basing condition, as one's own beliefs are at least the kind of thing with which one can be acquainted. I don't think a doxastic basing condition is satisfactory, however. First, take the Powerball example again, and add that I believe my belief that my ticket is a loser is based on my

---

<sup>79</sup> (Korcz 2015) is an excellent catalog and critique of different basing conditions that epistemologists have proposed.

examination of the odds. However, it remains true that my belief was caused by reading the horoscope. I think whatever the original intuition that basing is necessary had going for it equally tells in favor of a causal analysis over a doxastic analysis. Second, presumably only a *justified* belief about which reasons one's beliefs are based on will count toward *proper* basing: my belief B1 is *properly* based on C if and only if I have a *justified* belief B2 that C is my reason for B1. But if the analysis of doxastic justification makes reference to justified belief, the analysis will be circular.<sup>80</sup>

A general objection to the need for a basing condition is that once it is granted that the subject is aware of everything the proposition has going for it, there is nothing that could make it *more* probable for the subject that the proposition is true. My belief that my ticket is a loser does not become more probably true when, in addition to the fact that I am aware of the staggering odds that it is a loser, it is also true that the belief is *caused* by considering the odds rather than something else, or that it is also true that I believe the first-level belief resulted from considering the odds. My epistemic position with respect to the proposition that the ticket is a loser does not change in the slightest when a basing condition is met. Nor does it change when a basing condition is unmet.<sup>81</sup>

---

<sup>80</sup> This objection only goes through on a doxastic analysis of *proper* basing. It does not target doxastic analyses of basing *per se*. However, I cannot see the appeal of a doxastic basing condition that says B1 is based on C when one *unjustifiably* believes that C is the reason for B1.

<sup>81</sup> Lehrer (1971) makes a similar point with his case of the 'gypsy lawyer': the lawyer's confidence in the true proposition that his client is innocent is sustained by his evaluation of tarot cards, although the lawyer also has available evidence that his client is innocent. There is a complication in Lehrer's presentation, however: he claims that the lawyer's evidence yields knowledge that the client is guilty is because of the lawyer's unjustified faith in the cards. If that claim is just a description of the sequence of causes that led the lawyer to obtain his evidence, then I see no problem with the claim; but if it is instead a description of that in which the lawyer's basis for believing consists, then I think it undermines the earlier idea that the lawyer really is aware of the evidence that exonerates the client.

Many internalists take this sort of consideration to suggest that propositional justification is more epistemically significant than doxastic justification. I agree. However, I also think the very same reasons that support taking propositional justification to be more epistemically significant than doxastic justification support altogether rejecting proper basing as a condition on knowledge. If I find myself in a position where  $p$  is guaranteed to be true by factors I am aware of, it simply should not matter what else I might believe about why I believe  $p$ , nor should it matter by what route I came to affirm  $p$ . So long as I have a guarantee that  $p$  is true and I believe that  $p$  is true, nothing could improve my epistemic situation with respect to  $p$ . If nothing could improve my epistemic situation with respect to  $p$  other than my meeting the conditions that I do in fact meet, then nothing else is necessary for me to know that  $p$ . So, one can have knowledge when one's belief is not properly based on one's justification.

I suspect that part of the reason philosophers often think a basing condition is necessary for knowledge is that they have standards for knowledge that are lower than mine. Philosophers usually think one can know that a proposition is true when they meet standards that render the proposition probable to some degree less than I require, in addition to meeting some external conditions. In their case, adding a basing condition might help with connecting the alleged knower's belief to the truth of the proposition believed. Suppose that some analysis of knowledge says a subject knows a proposition is true when the subject believes the proposition, the proposition is true, the justification renders the belief overwhelmingly probable, plus some anti-Gettier condition. Then, on this analysis, one might have a belief that amounts to knowledge even though the belief's

connection to truth is really not that clear. The belief would be better connected to what makes it true if it were added that its justification—which *is* (presumably) connected to truth—is somehow *responsible* for the belief's existence. If something like that is the rationale for a basing condition, then internalist infallibilism does not require meeting any further non-epistemic conditions in order to secure a connection to truth, because on that view one only knows that *p* when one is directly acquainted with the factors that make *p* true.

That a weaker view of knowledge can give rise to worries about basing surfaces in Schaffer's presentation of his 'debasement demon' argument for skepticism (2010). He identifies two important requirements for knowledge that are distinct from truth, evidence, and belief: a non-accidental connection between one's evidence and the truth of one's belief, and that one's belief be based on one's evidence. Having specified these conditions, he imagines that there is a demon that prevents him from knowing what he takes himself to know by ensuring that his beliefs are based on reasons that do not justify them, even while he possesses good reasons for holding his beliefs. But on my view one need not worry about losing knowledge as a result of improper basing because the only kind of connection between evidence and truth I allow as appropriate for knowledge is necessarily non-accidental. When one is directly acquainted with the correspondence between one's evidence and the proposition that evidence makes true, there is no possible accident that could make it possible to satisfy that condition while the proposition is false. Since my analysis of knowledge needs no such non-accidentality condition, it also needs no basing condition.

The same line of thought suggests that a basing condition need not be met even for doxastic justification. Recall the Powerball example, where I have excellent justification for believing my ticket is a loser on account of the long odds, but the cause of my belief that my ticket is a loser was that I read a pessimistic horoscope. I am inclined to say, actually, that my belief that my ticket is a loser *is* justified. On that analysis of fallibly justified belief that I will shortly elaborate, one is fallibly justified in believing that *p* when one infers *p* from a known proposition and one is directly acquainted with a relation of making probable (to a degree less than entailment) that holds between the inferred proposition and the known proposition.<sup>82</sup> In that case, when one's belief is justified, one is again aware of everything the belief has going for it, epistemically speaking. Thus the same reasons for denying the need for a basing condition on knowledge also support rejecting a basing condition for doxastic justification. Since I have strict standards for justification as well as knowledge, I do not need a basing condition to help secure a connection between the justification for a belief and the truth of the belief. No further condition could be met that would make one's belief more likely to be true. As with knowledge, one can add a basing condition if one likes, but it will not put the putative knower or justified believer in a better epistemic position with respect to the proposition allegedly known or justifiedly believed. On my view, then, one has doxastic justification when one has propositional justification plus

---

<sup>82</sup> This condition is not actually sufficient because it faces two problems: the addition of fallibly justified beliefs further 'up' the inferential chain eventually results in beliefs that are unjustified, and the problem of defeat. I address these problems in the next section.

one believes the proposition for which one has propositional justification. No further connections are necessary.

It might seem that if no basing condition is necessary, there is at least a further condition much like a basing condition, and that condition is necessary. Suppose that while I do not need my belief that *p* to be based on its support in order to know or justifiably believe that *p*, I had better not believe that my reasons in support of *p* are bad. My awareness of the odds of losing justify believing that my Powerball ticket is a loser, but if I further *believe* that the odds are a bad reason to believe the ticket is a loser, perhaps that could destroy my doxastic justification or knowledge.

I am confident that if I believed my basis for believing the ticket is a loser was a defective basis, I would not continue to believe that the ticket is a loser. And if I did not believe the ticket is a loser, then I would fail to meet the belief condition of doxastic and justification and knowledge. But for the objection to really work against my argument that no conditions are necessary for doxastic justification or knowledge except for the conditions I proposed, the objection cannot amount to pointing out a way to destroy *belief*. Rather, we have to imagine that all of the conditions constitutive of doxastic justification or knowledge are met except for this further condition that I not believe my basis for believing is defective. So suppose I (somehow) believe both that my ticket is a loser on the basis of my awareness of the odds of losing, and I believe that the odds are a bad reason to hold the former belief. My belief that my evidence is bad is not capable of defeating the real evidential support my Powerball belief has if the evidence-belief is not itself justified. If it is just a stray doubt, it does nothing to undermine the actual support

had by my Powerball belief. So, only a *justified* belief that one's basis for believing p is defective is in the running as an additional constraint on doxastic justification or knowledge. In the Powerball case, I am at a loss as to what could justify a defeating evidence-belief while at the same time granting that I am directly acquainted with all of the conditions that constitute the justification for my Powerball belief and directly acquainted with the Powerball belief. I suggest that a lack of defeating evidence-beliefs is not an additional necessary condition on doxastic justification or knowledge; rather, the lack of a justified second-level belief that one's evidence fails to make probable one's first-level belief is implied by one's awareness of very strong grounds in support of the first-level belief. The condition that there be no justified defeating-evidence belief is redundant. I will further address the related topic of defeat for propositional justification in the next section.

There may be another explanation of the intuition that the Powerball belief caused by reading the horoscope is somehow epistemically defective even though it is justified. If the explanation is good it would help support my argument that proper basing is unnecessary. I now want to argue that there is something epistemically 'bad' in the scenario where the Powerball belief is caused by reading a horoscope even though the belief is supported by very good epistemic reasons: however, the 'badness' does not concern the belief's epistemic status, but the epistemic 'character' of an agent who trusts horoscopes. In ethics, there is a distinction between the evaluation of an agent's action and the evaluation of an agent. According to some versions of act utilitarianism, the right action is whatever action has the best actual consequences relative to other actions that

the agent could have performed in the circumstances. To modify an example from J.J.C. Smart, suppose someone near Berchtesgaden in the late 1800s saw a boy drowning in a river, and, being entirely without regard for human life, decided to simply finish his lunch while watching the boy drown.<sup>83</sup> Unbeknownst to this person, the boy was Hitler, so letting him drown had far better actual consequences than saving him would have had. Thus, the morally indifferent observer performed the right action, even while that person's motive is morally monstrous. Clearly only a bad person would simply choose to allow another person to suffer a painful death merely for the purpose of convenience, even if sometimes by luck acting from a bad motive results in a right action. I think theories of right action should allow for a distinction between the evaluation of actions and the evaluation of an agent's motives and character.

That distinction has a direct parallel in epistemology. We should distinguish the evaluation of an agent's beliefs and the evaluation of an agent's belief forming method on a particular occasion and an agent's overall epistemic 'character'. In the Powerball case, my belief is epistemically justified because it is supported by excellent epistemic reasons that make the belief highly likely to be true. However, my belief was formed by relying on a belief-forming method that I should not use. Plausibly, we should not trust horoscopes because the relation between what a horoscope 'says' will happen and what

---

<sup>83</sup> Smart (1973) did not invent the distinction between the evaluation of agents and the evaluation of actions, but he clearly articulates the distinction in the course of defending his preferred version of act utilitarianism. In his example from which mine is adapted, someone saves the drowning boy, who turns out to be Hitler. Smart uses the example to argue that insofar as right actions are optimific, saving the boy was wrong, even though the person who saved him by acting from the motive of saving life deserves praise for acting from that motive.

will happen is accidental.<sup>84</sup> So, we can say my belief is epistemically justified even though I should not hold it as a result of trusting my horoscope. Furthermore, we can say an agent has a defective epistemic ‘character’ if that agent regularly relies on belief-forming methods that are not conducive to forming epistemically justified beliefs.<sup>85</sup> What is bad about my epistemic character is that if acting from my epistemic character were to produce many true beliefs to false ones, it would be purely accidental. If I tend to hold beliefs on the basis of defective reasons, then I am a defective thinker.

The point of making all of these distinctions is this: the intuition that there is something epistemically amiss about me holding my belief that my Powerball ticket is a loser as a causal effect of trusting my horoscope is surely on to something, even while I have excellent epistemic reasons for holding just that belief. By making clear distinctions between different ways of evaluating beliefs, agents, and the methods agents use to form beliefs, there is room for argument as to what the intuition that there is something epistemically bad is tracking. It is surely the case that I should not hold my Powerball belief in just the way I do. However, I suggest that is because *I* am a blameworthy *believer* with respect to my Powerball belief. That is not a defect *in my belief*. My belief is epistemically very well-off. In this scenario, I have an epistemically justified belief

---

<sup>84</sup> Arguably, this grants too much: it does not take a careful reader to see that horoscopes predict virtually nothing specific.

<sup>85</sup> I say “not conducive” here, rather than “opposed to” forming epistemically justified beliefs. As I above characterized the problem with trusting horoscopes, the issue is that there is no positive correlation between trusting one’s horoscope and forming true beliefs as a result. We need not make the stronger claim that a belief-forming method should not be trusted only when that method tends to result in false beliefs.

that results in part from an epistemically defective character, much the way a person acting from a bad motive may still end up performing the right action.

### 3.2.3 *Foundational Knowledge, Inferential Justification, and Inferential Knowledge*

There has been some controversy among foundationalists as to the epistemic status of foundationally justified beliefs. Historically, foundationalists thought that foundationally justified beliefs are known, indubitable, certain, infallible, incorrigible: whatever the highest property of epistemic appraisal is, foundationalists used to think foundationally justified beliefs possessed that property. Foundationalism was later criticized on the grounds that there are no beliefs that possess the property foundationalists said foundationally justified beliefs have. Foundationalists replied by distinguishing foundationalism as a thesis about the *structure* of epistemic support from the *strength* of epistemic justification foundationally justified beliefs possess. Having made the distinction, foundationalists could keep their foundationalism while rejecting the more controversial claim that foundationally justified beliefs had exceedingly strong epistemic justification. Foundationalists came to accept so-called moderate foundationalism, which while characterized in different ways, roughly says that foundationally justified beliefs can be probable to a degree less than knowledge.

Propositional justification for some proposition  $p$  is foundational when that justification is not constituted by the justification of any other proposition. So, one has propositional foundational justification for a proposition  $p$  when one's justification for  $p$  does not include the justification for any  $q$  when  $q \neq p$ . One has a foundationally justified *belief* that  $p$  when one believes  $p$  and one has foundational propositional justification for

p. It is usually easier to talk about these issues in the language of doxastic justification (“justified belief that p”) rather than in more cumbersome terms of propositional justification (“there is justification for S to believe that p”), so I will do that. But ultimately, doxastic justification depends on propositional justification, but not the other way around.

One of the questions left unanswered at the end of the last chapter is how to distinguish justification and knowledge on a view that says the acquaintance relations that compose knowledge provide an infallible guarantee of truth. The problem is not that the concepts of knowledge and justification cannot be distinguished. It is that the analysis of knowledge does not obviously provide a place for epistemically justified beliefs that fall short of knowledge. I want to address that here.

First, I want to take a look at what some other recent epistemologists have said about the strength of justification for foundationally justified beliefs. There is no consensus here. Timothy Williamson (2000, p. 186) says “E=K suggests a very modest kind of foundationalism, on which all one’s knowledge serves as the foundation for all one’s justified beliefs. Perhaps we can understand how something could found belief only by thinking of it as knowledge.” Regardless of what he means by “modest,” it appears that Williamson defends a kind of strong foundationalism according to which a belief can only be epistemically justified to a degree less than knowledge when that belief is inferentially justified on the basis of one’s knowledge. Similarly, Timothy McGrew (1995) argues that foundationally justified beliefs are guaranteed to be true because their content is constituted by features of experience with which one is *de re* aware. On that

view, one cannot have a fallibly justified foundational belief. Furthermore, he defends an independent argument originally due to (Lewis 1946) that a belief would not be foundational if its epistemic probability were anything other than 1.0 (1995). Laurence Bonjour (2003, p. 64), when describing his view of foundational justification on the basis of experiences which include built-in awareness of content, says “[S]uch a non-apperceptive, constituent awareness of content might be said to be strictly *infallible* in something like the way foundationalist views have traditionally claimed (but which most have long since abandoned).” If foundations consist only in such beliefs, then Bonjour also denies the possibility of fallible foundational justification.

On the other hand, Chisholm seemed to allow for fallible foundations (Chisholm 1977, p. 24). He defines the epistemic property of being “directly evident”, his concept of foundational justification, this way: “h is directly evident for S =df. h is logically contingent; and there is an e such that (i) e is self-presenting for S and (ii) necessarily, whoever accepts e accepts h.”<sup>86</sup> He then defines “self-presenting” as follows: “h is self-presenting for S at t =df. h occurs at t; and necessarily, if h occurs at t, then h is evident for S at t.” Finally, he defines “evident” as one ‘step below’ certain (p. 12). So, putting it all together, the strength of justification for directly evident beliefs is just short of certain.

Robert Audi (1993, p. 106) defends a version of moderate foundationalism, by which he means:

---

<sup>86</sup> It is unclear exactly what Chisholm means by “whoever accepts e accepts h.” In this context, I think it is natural to understand the connection between e and h as a necessary connection that is discoverable ‘from the armchair’; in that case, Chisholm could have said ‘a normal person who is thinking carefully and who accepts e is thereby in a position to justifiably accept h’. But that is at least verbally quite different from “whoever accepts e accepts h.” So, I will not hang anything on my interpretation.

There are times at which S can have at least one justified corrigible belief whose justification does not, at these times, depend on, in the sense that it derives from, the justification of any other belief(s) S has at these times.

So, he allows fallibly justified beliefs whose justification could be defeated by other justified beliefs, so long as the justification of the former beliefs does not *derive from* the justification of any other beliefs.<sup>87</sup> Paul Moser (1991, p. 145) defends fallibly justified foundational beliefs. That is consistent with his view that justification is an explanatory notion. When a hypothesis h explains evidence e, it is often (usually) the case that h does not entail e. So, his view should (and does) allow for the possibility of fallibly justified foundational beliefs. Evan Fales (1996, p. 9 and ch. 6) argues that there are fallible foundational beliefs on the basis of the nature of given perceptual content. Richard Fumerton (2006b, 2010, 1995) argues that fallible foundational beliefs are consistent with his acquaintance theory of noninferential justification, for a few reasons: he recognizes degrees of correspondence, so when one is acquainted with the thought p and the fact q (not p) and a high degree of partial correspondence between them, one can get foundational justification even for a false belief; and because he recognizes degrees of similarity between facts (also implicit in the previous remark about partial correspondence).

---

<sup>87</sup> Audi's quote states that there can be foundationally justified "corrigible" beliefs. In his paper, "corrigible" is used to mean just what I mean by "fallible": he calls an *incorrigible* belief one that, when held, amounts to knowledge (p. 98). It follows that a belief that is corrigible is one that may be held without amounting to knowledge. That, combined with the claim that the belief that is not known is justified and foundational, and it becomes clear that what he means by 'corrigible foundational belief' is what I mean by 'fallible foundational belief'.

The survey of epistemologists just given shows that there are two central issues at play in the debate about the strength of noninferentially justified beliefs. First, there is the issue of whether or not foundationalism as such entails that foundationally justified beliefs must be certain knowledge. (McGrew seems to think so; the others I have mentioned are opposed.) Second, there is the issue of whether or not the correct analysis of noninferential justification is consistent with fallible noninferential justification. (Fales, Fumerton, and Moser answer affirmatively, Bonjour, negatively.) It is important to distinguish these two issues because one gets different versions of foundationalism and thus the possibility of different objections to one's favored version of foundationalism depending on the view one takes on each issue. For instance, if one could show that there are no infallible foundational beliefs, one would have thereby committed McGrew to rejecting foundationalism: because he takes foundationalism to imply strong foundations, it follows that if there are none, foundationalism is false. However, one could also take the view that foundationalism *qua* foundationalism does not entail infallible foundations, but that the best way of making sense of foundational justification is that only infallibly justified beliefs are foundational. On that view, while foundationally justified beliefs do not *have to be* infallible on pain of rejecting foundationalism (*pace* McGrew), they just are (*pace* Fales, Fumerton, and Moser).

While I do not think internalist infallibilism *entails* that there are no fallible foundations, I do think the view fits best with the further commitment that all foundational justification is (infallible) knowledge. I reject arguments that say foundationalism or even internalist foundationalism requires infallible foundations.

However, I think the best way to make sense of plausible examples of foundationally justified beliefs is that such beliefs amount to knowledge, according to the definition of knowledge I have been defending.

Why think that? Well, the motivation for strong standards for knowledge is the thought that we are in a position to decisively answer questions about what we know. That suggested the view that knowing that  $p$  consists in being directly acquainted with the thought  $p$ , the fact  $p$ , and the correspondence between them. When those conditions are met, one's belief that  $p$  is guaranteed to be true. I think the concepts of justification and knowledge share their epistemic conditions, but justification is something weaker than knowledge. Justification is knowledge-lite. But having defined knowledge the way I have, the question of fallible justification comes to this: how can I dilute knowledge so as to yield knowledge-lite?

One way to get something less than knowledge while still relying fundamentally on direct acquaintance with the factors that make one's belief true is to hold that fallibly justified beliefs are always based on one's knowledge. I suggested this view in the previous section. One could fallibly justifiably believe that  $p$  on the basis of one's knowledge that  $q$  when one is acquainted with everything that constitutes knowledge of  $q$  and be acquainted with a relation of making probable that holds between  $q$  and  $p$ , while one is also acquainted with one's belief that  $p$ . On that view, fallible justification arises from being directly acquainted with a support relation that is weaker than entailment, and

so inferring from one's knowledge is not guaranteed to result in a true inferred belief.<sup>88</sup> Since one could be acquainted with all of those factors and yet have a false belief, one's justification for the belief that p is fallible.

But how to account for fallible *foundational* justification is altogether more complicated. Whether or not one should recognize fallible foundational justification ultimately depends on the plausibility of a proposed account of fallible foundational justification that is consistent with an acquaintance-based theory of justification and knowledge. I will now consider some challenges to such a theory. In my view, it is not clear that an acquaintance-based account of fallible foundational justification *cannot* be given. However, I think the challenges to such attempts are forceful enough that when viewed alongside the plausibility of infallible foundations, the scales ultimately tip in favor of recognizing only infallible foundations. If that is correct, then all fallible justification is inferential.

Fumerton has argued (2010, 1995) that on his view foundational justification for a false belief is possible, and would work like this: S is acquainted with the thought p, a fact *similar to* and easily confused with p (call it q), and a relation *similar to* correspondence holding between p and q. Whether or not that account is acceptable depends both on whether adequate analyses of fact-similarity and of relations similar to correspondence can be given. The notion of fact-similarity is intuitive. There are some sensations that I might classify as either a pain or an itch, but which I surely would not

---

<sup>88</sup> Keynes (1921) argues that there are logical partial implication relations that hold between propositions. I follow Fumerton (1995) who makes use of such relations in providing an account of inferential justification.

classify as the taste of chocolate ice cream. So, it seems that some facts are highly similar to others, enough so that one might be unable to distinguish which sensation one is experiencing. While that only suggests a direction for further analysis, the plausibility of that kind of example is all we should need to make sense of fact-similarity for the purpose of providing an account of foundational justification.

I am more worried about the suggestion that acquaintance with a relation ‘similar to correspondence’ helps provide justification for believing *p* when one is acquainted with that relation holding between *p* and a fact similar to *p* (*q*). The relation in question is a kind of correspondence that is weaker than what we might call ‘exact’ correspondence. While correspondence is not a kind of resemblance, both are relations of ‘matching’ or ‘fit’: so, comparisons with ways in which objects can resemble or represent one another should be helpful for getting a conceptual grip on the concept of correspondence. For example, some pictures correspond better to the objects they are representations of than other pictures do. A drawing my preschooler son makes of a giraffe corresponds to giraffes *better* than a drawing my son makes of a cup does, but my drawing of a giraffe corresponds to giraffes better than my son’s drawing of a giraffe does. Neither of us comes anywhere near a (normal, good) photo of a giraffe. Using the same notion of similarity, one might say that correspondence comes in degrees, with exact correspondence as a limit, and lesser degrees of correspondence obtaining between two corresponding entities as their similarities erode further down the scale.

It is overwhelmingly natural to think of some representations as more/less similar to the objects they represent than other representations are. (Again, I am not assuming

that correspondence is a kind of resemblance.) However, I don't think that is enough to support the view that the relation that obtains between distinct entities is a relation of partial resemblance analogous to partial correspondence. Take the example of a picture, again. Suppose there are two very similar but visibly different pictures. The natural thing to say is that they look similar. But why? One thing that could be said is that the two pictures stand in a relation of partial resemblance that cannot be further analyzed. But when looking at these two pictures that are visibly distinct we would be able to point to the differences that prevent them from being identical, from corresponding exactly. Some of the other features, though, are the same. We can look at the two pictures and isolate similarities and differences. So while it is natural to say they 'look similar but not identical,' thinking more about why we make that judgment suggests that the relation the two pictures bear to one another is not an unanalyzable relation of similarity, but a relation of exact resemblance between many of the pictures' features but not all of them.<sup>89</sup>

If that is right—if partial resemblance is to be understood in terms of exact resemblance between parts in a complex—then it seems plausible to expect partial correspondence to have similar features. First, the motivation has been undermined for claiming that there is a *sui generis* relation of partial correspondence. Second, partial

---

<sup>89</sup> Fodor and LePore (1992) deploy a similar argument in support of their claim that psychological holism has damaging consequences for psychological explanation. Their argument is that the only robust notion of similarity is partial identity. So, if psychological holism is true, then two thinkers with different holistic content can be said to have similar thoughts only if their thoughts are partially identical. But since holism says an atomic thought gets its content from the whole system, there are no two token thoughts of distinct cognitive systems that are type-identical. So, they argue, holism is not so innocuous as many philosophers have thought.

correspondence will obtain only between complexes. The sort of example that motivated this discussion is most plausibly an example of mental properties that are simple.

Suppose I have a sensation that may be an itch or a pain, but I'm sure is not the taste of chocolate ice cream. If my thought that I am in pain partially corresponds to the sensation I am acquainted with in the sense that some constituents of my thought exactly correspond with some of the constituents of my sensation, that requires my sensation to have constituents. I have no reason besides this argument to think pain sensations literally have constituents.

However, some examples of inexact resemblance do seem to hold between simples. Consider this case. Before me are three color patches of the same shape, the same distance from me, and etc.—they share every feature except for color. Two of the patches are shades of red, the third is a patch of black. The two reds are the closest they can be in the color spectrum without becoming indistinguishable. It seems pretty obvious that the two red patches are very similar to each other while remaining non-identical, and that each is more similar to the other than either is to the black patch. Furthermore, the features of the color patches being compared are simple, for all that is being compared is the patches' colors. So perhaps we have located a relation of weak resemblance that cannot be analyzed as exact resemblance between parts of complexes. Perhaps that notion of weak resemblance could be put to work in understanding a notion of weak correspondence that can be used in an account of fallible foundational justification.

As I said, my goal is to argue that all fallible justification is inferential on the basis that it is difficult to make sense of how fallible foundational justification is possible

and easy to make sense of infallible inferential justification, not because fallible foundational justification is unintelligible or because foundationalism entails that foundational justification is knowledge. So, now that I recognize both a natural way of analyzing partial correspondence as exact correspondence between parts of complexes, and the plausibility of unanalyzable partial correspondence between simples, the relevant question is this: can the notion of unanalyzable partial correspondence that seems present in the case of very similar color patches be effectively used to provide an account of fallible foundational justification?

Suppose I am seeing patches of red #111 and red #112, but not red #113. Now, suppose I falsely believe I am seeing red #113, and (for the sake of argument) that I am acquainted with my belief that I am seeing red #113, the facts that I am seeing red #111 and red #112, and the unanalyzable partial correspondence between my belief and the fact that I am seeing those two reds #111 and #112. Could that provide me with noninferential justification for the false belief that I am seeing red #113? I frankly have a hard time imagining what it would be like to have so fine-grained a belief about colors. But remember, the plausibility of an account of fallible foundational justification depends on the items confused being very similar to one another.

Constructing a more promising example of fallible foundational justification will require coming up with a thought that one can be acquainted with, but which still allows for the same narrow possibility of mistake as that of judging that some color patch is such-and-such rather than so-and-so. Here is an example Fumerton has used for roughly this purpose (2010, p. 380-1). Suppose pain states occur along a continuum that

eventually drops off into non-pains. Each state along the continuum is indistinguishable from the one next to it, just as the color patches in the earlier example were indistinguishable: there might be a number of pain states 1-1000, and I could be in pain #1000 that is indistinguishable from non-pain #1001. (What would state #1001 be? An itch?) I might have pain that is right up against the ‘border’ of non-pain. Even if I were acquainted with such a pain, I would not be able to tell whether the state I am acquainted with is a pain versus a non-pain one ‘stage’ away. As a result, my belief that I am in pain is very ‘nearly’ false. If that is all plausible, then it also seems plausible for me to be acquainted with the thought that I am in pain, a state just on the *other* (non-pain) side of the ‘border’, and the partial correspondence between them. Fumerton suggests that in such a scenario, one could have foundational justification for the false proposition that one is in pain.

It’s not clear to me that in this scenario it should be granted that it is probable *for me* that I am in pain rather than not. After all, per hypothesis, my state is *indistinguishable* from pain. I am unable to tell whether I am experiencing pain that is nearly non-pain, or non-pain that is nearly pain.<sup>90</sup> I am acquainted with some vaguely uncomfortable mental state, but that’s the most precise characterization I can give. There is nothing in my perspective that tells in favor of the experience being a pain experience

---

<sup>90</sup> In a new paper, Chris Tucker argues that the acquaintance theory of noninferential justification cannot accommodate fallible noninferential justification (2016). If I read him correctly, he and I are relying on the same crucial assumption, which he states in section five as follows: “[A]wareness of X’s similarity to things of type T provides justification for thinking X is T only if one has some antecedent justification for believing that X’s having these similarities reliably indicates being of type T.” He goes on to insist that the antecedent justification must be empirical, but I see no basis for that. The argument succeeds if it requires any kind of antecedent justification, including *a priori* justification.

rather than not. It seems that if there is any justification for believing that one is in pain when one is so nearly not in pain, it is because that justification is inferential.

Two points deserve stress here. First, in this scenario, what is probable for me is that I am in a state *much like* pain, that I am having an uncomfortable experience. But my justification for believing that, while foundational, is also infallible. I am directly acquainted with my state that is much like pain (whether it is pain or not), the thought (so long as I'm thinking it) that I am in this uncomfortable state, and the exact correspondence between them. Because that straightforward account of the case can be given, I see little incentive to grant that one can have justification for believing that one is in pain when one is not. But if one continues to feel the pull of the intuition that one is justified in believing that one is in pain in this scenario—this brings me to my second point—one could always account for it this way: I know that I am in a state much like pain; in addition, I am acquainted with the thought that I am in pain, my knowledge that I am in a state much like pain, and the relation of making probable holding between my knowledge that I am in a state much like pain and the thought that I am in pain; therefore, I am in a state much like pain. My justification for believing that I am in pain would then be fallible, but it would also be inferential.

The same account can be given of other cases that initially suggest there are fallible foundations. Evan Fales (1996, p. 141) presents a case in which he lived on a farm where there were feral cats, including a black one and a grey one. He reports that he was “for several reasons concerned to take notice” of the “comings and goings” of the

cats<sup>91</sup> and once had a visual impression of a running cat that was definitely, determinately either black or grey, but he could not tell which. He takes this case to suggest that a probability of less than 1 can be assigned on the basis of this experience to the proposition that the cat-experience was black, and that a probability of less than 1 can be assigned to the proposition that the cat-experience was grey. In keeping with what I said about Fumerton's pain experience example, I suggest the alternative construal that the fallible justification that the cat-experience was some determinate color C is not based solely on the experience, but on the foundational *knowledge* that he had a black-or-grey experience, from which the proposition *the cat-experience was black* and *the cat-experience was grey* can both be inferred.

There is a final reason for thinking I cannot have fallible foundational justification. When we hope to be justified in holding some belief, we are hoping that the belief we hold is likely to be true. Ideally, then, an analysis of justification will have a conceptual connection with truth, even if justified beliefs usually turn out to be false. On my analysis of fallible justification, it is analytic that when one is fallibly justified one's belief is likely to be true. That is because I analyze fallible justification as supported by logical relations to one's knowledge, which analytically includes truth. Fallible foundations do not allow for such a conceptual connection between justification and truth.

It is often thought that the concern for connecting justification to truth favors externalism. One version of that concern applies to me: on my view it may be that one's

---

<sup>91</sup> Why, Evan?

belief is in a sense defined above likely to be true even though it is *objectively* unlikely to be true because one's beliefs may be produced by unreliable processes or operating in the wrong environment. My analysis of justification will not satisfy the epistemologist who wants an analysis of justification to call a belief 'likely to be true' only if beliefs of that type are often true. But, as Cohen (1984) argued with his 'new evil demon' thought experiment, two mentally identical counterparts, one of which is unlucky enough to inhabit a world where beliefs formed on the basis of experience are nearly always false, are intuitively equally justified. So the problem I am pressing against fallible foundations is not that they allow for beliefs to be justified when those beliefs are objectively unlikely to be true, because I reject as a necessary condition on justification that justified beliefs be objectively likely to be true. My argument is that only with infallible foundations is there an analytic connection to truth. It is analytic, on my view, that when a belief is justified it is likely to be true, because justification is (partly) defined in terms of its relation to knowledge which, in turn, is (partly) defined in terms of truth. So, an analysis of justification that allows only infallible foundations easily respects the idea that justified beliefs are likely to be true. I do not see how fallible foundations can accommodate that understanding of justification.

So, while I think the case for infallible foundations is not conclusive, it has a lot going for it: enough to ultimately weigh in favor of infallible foundations and the corollary that all fallible justification is inferential. To summarize, the reasons were that it is far easier to provide an analysis of infallible foundational justification than of fallible foundational justification given the technical problems associated with defining partial

correspondence and fact-similarity; that the borderline cases that might seem to lend support to the thesis that there are fallible foundations are easier to account for on the view that justification in those cases is inferred from infallible foundations; and because only infallible foundations secure an analytic connection to truth.

To be explicit, this is the analysis of justification I propose. I start with the analysis of foundational knowledge from the previous chapter:

S knows that p =df. (1) S is acquainted with the belief that p  
(2) S is acquainted with the fact p  
(3) S is acquainted with the correspondence between the belief that p and the fact p.

Knowledge is the limit case of justification, so trivially what one knows one has justification for believing. This is a partial analysis of *fallible* propositional justification.

Below, I explain why it is only partial.

S is fallibly justified in believing that p =df. (1) S knows that q (q≠p)  
(2) S is acquainted with a logical relation of making probable holding from q to p  
(3) S is acquainted with the thought that p

For fallible doxastic justification, just add belief. Finally, I have not said anything about the possibility of inferential knowledge. Because knowledge requires being in a position to eliminate any possibility of error, it is not possible to acquire knowledge from inference through fallibly justified belief. So, whatever is known by inference is based on other knowledge. However, given the strict standards for knowledge, inferential knowledge consists only in further acts of acquaintance with thoughts and the facts that guarantee their truth.

S knows that p on the basis of q =df.

- (1) S knows that q ( $q \neq p$ )
- (2) S is acquainted with the *entailment* from q to p
- (3) S is acquainted with the thought that p

Only entailment preserves truth from q to p, so only entailment can serve as an acceptable logical relation between foundational knowledge and further knowledge that is inferred on its basis.

To put all of this in a way that makes more obvious the fact that the analyses of justification, foundational knowledge, and inferential knowledge all have the same basis, here are those definitions again. First, the analysis of knowledge:

*S knows that p* if and only if (1) S is acquainted with the fact p, the thought that p, and the correspondence between the fact p and the thought that p...

OR (1) S is acquainted with the thought that q ( $q \neq p$ ), the fact q, and the correspondence between the fact q and the thought that q, AND (2) S is acquainted with the thought that p and the entailment from q to p...

OR (1) S is acquainted with the thought that r ( $r \neq q \neq p$ ), the fact r, and the correspondence between the fact r and the thought that r, AND (2) S is acquainted with the thought q and the entailment from r to q, AND (3) S is acquainted with the thought p and the entailment from q to p...and so on.

To make clear that fallible justification has the same basis as knowledge, here is a *partial* analysis of fallible justification.

*S is fallibly justified in believing that p* only if (1) S is acquainted with the fact q, the thought that q, and the correspondence between the fact q and the thought that q, AND (2) S is acquainted with the thought that p ( $p \neq q$ ) and the making probable relation holding from q to p...

OR (1) S is acquainted with the fact r, the thought that r, and the correspondence between the fact r and the thought that r, AND (2) S is acquainted with the thought that q ( $q \neq r$ ) and the making probable

relation holding from r to q, AND (3) S is acquainted with the thought that p ( $p \neq q$ ) and the making probable relation holding from q to p...and so on.

As these statements of the analyses of fallible justification and inferential knowledge show, on my view the fundamental epistemic concept is knowledge, which itself is analyzed in terms of acts of acquaintance with the facts that make a belief true.

I said that the analysis of fallible justification is partial and offered only necessary conditions. There are two reasons for this. First, one might meet the conditions specified for fallible justification for p above while it is true that if one had thought about things differently, one might have used another available inference from a different bit of one's knowledge and one's acquaintance with a making probable relation between that knowledge and another proposition that entails that p is false. In other words, what would otherwise be an acceptable route to justification can sometimes be defeated. So, the conditions for fallible justification given above are insufficient. I will avoid entering into a further controversy on how to give a satisfactory sufficient condition for justification. I am content to point out that while the analysis is still consistent with the claim that we are in a position to decisively answer questions about what we are justified in believing, it will often be more difficult to discover all of the grounds that constitute one's being fallibly justified in believing that p on account of the additional reflection necessary to ensure that one does not have additional evidence that undermines one's other grounds for p. Many epistemologists take this problem to motivate a so-called total evidence requirement on justification, which says that one has justification for believing a proposition only when that proposition is supported by *all* of one's evidence, and not merely whatever part of one's evidence one is focused on. I support that move. But in

addition to further evidence, part of the problem is supposed to be that there are other legitimate inferences one could have performed in order to reach a different conclusion on equally well-supported grounds. So, providing a sufficient condition for fallible justification also involves some statement to the effect that the proposition is supported by all or some sufficiently weighty subset of the inferences one could have performed on one's total evidence.<sup>92</sup>

Second, there is a problem of accumulating fallibility. Plausibly, there is some threshold beyond which one fails to have justification for believing a proposition, even if one's belief is supported by acquaintance with making probable all the way down to a basis in knowledge. When one makes fallible inferences, the inferred proposition is less probably true than the proposition from which it was inferred. Do this enough and the result is a proposition that is unlikely to be true on one's evidence. (For an easy illustration, just keep multiplying .95 over and over and watch the product decrease each time.) This problem, like the first one, points to the failure of the analysis above to provide a sufficient condition. And, like the first problem, I do not have an original solution to offer. The moral here is that one can't get too far from one's foundational knowledge and still have reasonable confidence that one's beliefs are likely to be true. If one met the necessary conditions stated above in the analysis of fallible justification, and

---

<sup>92</sup> Read literally, evidentialism in epistemology suggests that such a latter condition is unnecessary, because evidentialists insist that justification supervenes on evidence and evidence alone (Feldman and Conee 1985). But justification requires more, including awareness of connections between one's evidence and the propositions supported by the evidence, and perhaps more than that still: coherence, proper basing, etc. So, evidentialism as understood in the recent evidentialist tradition is false.

one's justification was not defeated and was below the threshold for justification, then one would be justified in believing the proposition.

### 3.3 The Extent of Knowledge and Justification

Much of the most difficult epistemological work is done before reaching the point where the question of the extent of our knowledge and justification is asked. Given what I've already argued, the extent of knowledge can be sketched easily: when one is acquainted with a belief and its truth, then that belief is knowledge; otherwise, it isn't, although one might still be justified in holding the belief. Whatever analysis of a concept is offered determines an extension, so questions about what is in the extension are secondary to the nature of what it is that determines the extension.<sup>93</sup> But if there are problems that arise related to the proposed extension—say, the proposed extension is intuitively much smaller than what we reflectively think the extension is, or some things are included in the proposed extension that seem like they should not be while others are excluded that seem like they should be—then that can be used to motivate an objection to the analysis provided. Although the next chapter discusses methodological concerns in more detail, I want to explicitly address the extent of knowledge here for the reasons just given.

---

<sup>93</sup> Strictly speaking, probably no one has ever committed to the view that 'intension determines extension,' for once an intension is provided, the extension still depends in part on whether or not there are things in the world that fit the intension. Even provided an intension of 'tiger,' whether or not "x is a tiger" is true still depends on whether or not x has the properties stated in the intension of 'tiger'. Still, there is a clear difference between the sort of view that recognizes the conceptual priority of a thinker's intension (or 'sense' or 'concept' or 'analysis') over extension, and the sort of view that attempts to take the former out of the picture altogether.

### 3.3.1 *Knowledge of the External World*

In order to find a secure bit of knowledge to serve as a foundation for a wholly legitimate reconstructed system of beliefs, Descartes wanted to find some beliefs that could not be called into doubt. The way he set about to accomplish this in the *Meditations* was to imagine a series of increasingly severe skeptical scenarios that implied that nearly all of the beliefs he held could be false. He eventually arrived at one certain belief: that whenever he is thinking, he exists (in some form or other). He discovered his secure foundations in the mind. Similarly, Russell starts his *The Problems of Philosophy* (1912) by distinguishing the way physical objects appear to us from the way we believe they are. Here he discusses the challenge of accounting for the physical object behind the appearance of a brown, wooden table:

To make our difficulties plain, let us concentrate attention on the table. To the eye it is oblong, brown and shiny, to the touch it is smooth and cool and hard; when I tap it, it gives out a wooden sound. Any one else who sees and feels and hears the table will agree with this description, so that it might seem as if no difficulty would arise; but as soon as we try to be more precise our troubles begin. Although I believe that the table is 'really' of the same colour all over, the parts that reflect the light look much brighter than the other parts, and some parts look white because of reflected light. I know that, if I move, the parts that reflect the light will be different, so that the apparent distribution of colours on the table will change. It follows that if several people are looking at the table at the same moment, no two of them will see exactly the same distribution of colours, because no two can see it from exactly the same point of view, and any change in the point of view makes some change in the way the light is reflected (pp. 8-9).

Russell identifies changes in properties he is aware of that are not plausibly changes in the table's properties. Here he focuses on changes in colors. But he connects this issue of perspectival variation of color with the more general phenomenon of accounting for

differences in how things appear and how they are. His example is that of the painter, who has to attend carefully to the way things appear if he is to have any hope of making a realistic-looking painting of a table. If the table really is oblong, then as one walked around the table it would appear to shift from elliptical, to a sort of roundish-figure with one nearly flat side, to a triangle-like shape with rounded ‘corners’, and back to elliptical upon making it to the other side. We are so accustomed to forming beliefs about how things are on the basis of how they appear that we ordinarily fail to notice that things do not ‘look’ the way we think they really are.

What both Descartes’ methodical doubt and Russell’s observations about appearances suggest is that our foundational empirical knowledge consists in propositions about subjective appearances. In the case of Russell’s table, what the argument is supposed to show (and I am convinced that the argument does show it) is that what we are directly acquainted with in perception are not the physical objects we ordinarily take ourselves to be aware of, but something else which serves to represent the object (wholly or partially).<sup>94</sup> While what we are acquainted with changes, presumably the table itself does not undergo a change in its intrinsic properties, nor should the observed variation be accounted for by claiming the table somehow alters its projected appearance to accommodate observers’ changing points-of-view. The epistemological upshot is that whatever justification or knowledge we have for believing propositions about physical

---

<sup>94</sup> I will remain silent on metaphysical issues concerning the nature of such representation and the analysis that should be given of whatever allegedly represents the mind-independent object. Russell uses the discussion to motivate a sense data theory of perception, but I think it is possible to give an adverbialist analysis that equally well accounts for perspectival variation, illusion, and hallucination.

objects on the basis of perception is inferred from foundational knowledge we have about appearances. That is because what the argument shows is that we are not acquainted with physical objects, but with their appearances (however that is to be understood). Since all knowledge and justification reduces to facts concerning what we are acquainted with, all the knowledge and justification we could have on the basis of perception depends on foundational knowledge of appearances. So if I am justified in believing that the table before me is brown and oblong, at least this much is true: I know that I seem to see something brown and oblong and I am acquainted with some relation of epistemic support that holds between what I know about the table-appearance and the proposition that the table before me is brown and oblong.

The first major obstacle to providing an analysis of knowledge that accommodates the thought that we have lots of ordinary knowledge is how to get from knowledge of one's experiences to knowledge of the external world. Given the strict standards for knowledge I accept, there is a major hurdle that appears as soon as it is recognized that empirical knowledge depends upon knowledge of one's experiences: to know anything about the external world requires being aware of *entailing* epistemic relations that say that if some particular known proposition about one's experience is true, then some particular proposition about the external world is true. It is doubtful that we are aware of any such relations because it is doubtful that such relations exist, at least any that would be useful in getting us from knowledge of experiences to knowledge of the external world. The sort of relation we are looking for would get one from knowing that one is having an experience as of a brown, oblong thing to knowing that there is a

brown, oblong thing. So the entailment principle would relate having a brown, oblong experience and there being a brown, oblong thing, and from that inference plus other justified background assumptions to the conclusion that there is a brown, oblong table. But surely it is possible that there is a brown, oblong table-experience and no brown, oblong table: if one had the experience but was hallucinating, or the victim of some sort of deception, one could have exactly that experience when there is no physical object of the requisite sort responsible for the experience.

If we want to account for being justified in believing on the basis of sensory experience that physical objects exist, then instead of requiring acquaintance with a principle that says if one has such-and-such an experience, then such-and-such a physical object exists, one must be acquainted with a principle that says such-and-such an experience makes it probable that such-and-such an object exists.<sup>95</sup> While such principles cannot be so easily undermined by pointing to the possibility that the antecedent conditions hold while the consequent conditions do not—after all, they only say that something is *probable* given that something else is the case—it is not obvious either that there are such principles that can be used to acquire justified belief in an external world or that we are acquainted with those principles. To be clear, I am not saying none exist. There is good reason to think there are some necessary principles of nondeductive reasoning that we are aware of. It seems obvious that one is inferring well

---

<sup>95</sup> Fumerton (1995) argues for further restrictions on such principles and our awareness of them, and concludes that it is not clear that we can use such principles to avoid skepticism.

when one infers from (1) 99% of F are G and (2) x is an F to (3) x is a G.<sup>96</sup> But it is less obvious that there are necessary principles of nondeductive reasoning that can help get us justification for believing propositions about the external world. Is it really the case that if I am appeared to as though there is a brown, oblong table then it is probable that there is a brown, oblong table, and that I can be directly acquainted with this principle? While I find it rather obvious that the principle of statistical reasoning just cited is a good one, it is far from obvious to me that the epistemic principle used to infer the existence of a table is also legitimate. If it is legitimate, we may have lots of ordinary justification. But if it is not, the skeptical implications are rather severe.

### 3.4 Chapter Conclusion: Skeptical Epistemology

Internalist infallibilism plus the claim that we are directly acquainted with the contents of experience and not with physical objects in a mind-independent world implies that we know very little of what we say we know, and we are possibly not justified in believing much of what we take ourselves to be justified in believing.<sup>97</sup> Understood one way, the view does not imply that knowledge of mind-independent truths is *impossible*. It only

---

<sup>96</sup> Other conditions are necessary. As Nelson Goodman (1955) argued, the properties F and G must be ‘projectible’ properties: while the analysis of projectibility is controversial, it is uncontroversial that properties whose instantiation is indexed to times are not projectible. Goodman’s predicate ‘grue’ was analyzed as ‘green before time t and blue after time t’. It would appear that we could infer from ‘all observed emeralds are green’ and ‘this is an emerald’ to *both* ‘the next observed emerald will be green’ *and* ‘the next observed emerald will be grue.’ However, these predictions are incompatible, for if the next observed emerald is observed after t, the green-prediction says the emerald will be green while the grue-prediction says the emerald will be blue. The upshot is that there are restrictions on the properties (or predicate expressions) that can be used in nondeductive inference.

<sup>97</sup> Bertrand Russell thought he was acquainted with Platonic universals in addition to sense-data. I don’t think I am acquainted with universals. But I also don’t want to make acceptance of my overall view about the extent of knowledge depend on rejecting acquaintance with universals. Even if we are acquainted with universals, it will not help overcome the problem of skepticism described above.

says the extent of knowledge is more limited than many epistemologists think. If we had been beings endowed with direct acquaintance with facts in an external world, perhaps we could have some knowledge of an external world, other minds, the past, the future, and so on. Understood another way, to imagine beings with such capacities is to imagine beings very much unlike us, so it is necessary that any beings like us lack knowledge of anything but their own minds and easily understood necessary truths.

So, the view I have been defending implies a kind of local knowledge skepticism rather than global knowledge skepticism: we can have some knowledge, so global knowledge skepticism is false. Having knowledge analytically entails having justification, so by granting that we have some knowledge I have thereby granted that we have some justification. But perhaps not much. It depends on whether or not we are aware of principles of nondeductive inference that can connect our limited knowledge with propositions about mind-independent truths.

It is natural to object at this point that the epistemological work done here has surely been for nothing: the epistemologist's goal is to provide an analysis of the ordinary knowledge that we in fact have. One might assume that a key methodological presupposition is that our beliefs have the epistemic status we ordinarily take them to have: that we have, as we usually assume, an abundance of knowledge and even more highly justified beliefs. If the result of epistemological inquiry implies that this methodological presupposition is false, then one must conclude that the inquiry went wrong somewhere: I got the wrong answer and just have to start over. However, I don't accept that methodological presupposition. As I argued in chapter one, the interests of

the philosopher doing conceptual analysis should be primarily on figuring out the contents of concepts of philosophical interest. It is compatible with that endeavor that we never use those concepts to think true thoughts.

But I feel the pull to have ordinary ways of thinking turn out 'okay' post-analysis. At this point I could concede a limited kind of externalism: while externalist analyses of knowledge and justification are not strictly adequate in the sense that they do not adequately describe our concepts of knowledge and justification, they play a (perhaps very) useful role in accounting for much of our ordinary thought and talk. I want to continue to say that I know where my car is parked, where I'll be in May, who the President is, that I have hands, and so on. Well, then I could say that externalism offers a useful kind of revisionism: 'forget about the concepts of knowledge and justification you started with! They're useless, they won't help you find your car or plan your vacation, they won't help you classify others into reliable informants and unreliable informants...etc.' Externalism could offer a kind of revisionism that lets us continue to deploy epistemic concepts in ordinary judgments the way we are used to doing.

I could say that, but I won't. There is another way to have much ordinary thought and talk turn out 'okay'. That other way is by providing a defense of the idea that our ordinary thought and talk involves precisely those high epistemic standards I've been trying to articulate. On the revisionist externalist view, we reinterpret our epistemic talk in a way that *substitutes* an externalist analysis of epistemic concepts for the genuine article. On the view I propose, we reflectively accept that our epistemic talk is nearly always false. It may seem surprising to suggest that our ordinary epistemic concepts set

such high epistemic standards given that we make epistemic judgments using those concepts all the time. I will argue in the next chapter that we actually do talk about knowledge in a way that positions us to recognize that such talk is regularly false. I will defend a loose talk account of the semantics of knowledge claims against competing accounts like epistemic contextualism and moderate invariantism. The loose talk account says our knowledge talk is loose in the same way that much of our ordinary talk is loose: it is imprecise, and thus false, but we are often implicitly aware of the fact that such talk is false. As such, the argument of the next chapter both provides an error theory that helps support the theory of knowledge defended thus far by showing that it is compatible with common practice, and it provides an independent argument for demanding standards for knowledge by showing that our ordinary practice itself suggests high standards for knowledge.

## **CHAPTER FOUR:**

### **SKEPTICISM AND KNOWLEDGE ATTRIBUTIONS**

#### *Chapter Preview*

Anyone who accepts the skeptical epistemology defended up to this point has questions to answer about the implications of that skeptical epistemology. In this chapter, I focus on perhaps the most important question: why do speakers so regularly speak falsely when attributing knowledge to themselves and others? The question concerns the error theory skepticism needs. The issue of why knowledge attributions that are usually false are made so regularly by speakers is particularly important for someone who accepts views about philosophical analysis as I do. For I hold that philosophical analysis is an attempt to clarify for oneself the meaning of concepts of philosophical importance. And since concepts are often expressed through language, this method of philosophical analysis applied to the analysis of knowledge consists partly in examining the conditions under which it would be true to say “S knows that p” or “S does not know that p.” Here is the problem: there are a variety of knowledge attributions that seem natural to make, but if they are true, then skepticism is false. So the skeptic has to explain why knowledge attributions that seem natural to make are false, and the explanation will be far easier to accept to the degree that the skeptic succeeds in explaining how those knowledge attributions are somehow appropriate while false.

As we will see, the same metaphilosophical views that make the problem of competing accounts of “S knows that p” particularly acute for me also point the way to solving the problem. Philosophical analysis requires semantic competence; otherwise, it

would be impossible to correctly describe the concept under consideration. Philosophical analysis presupposes that the analyst is competent with the concept or term under scrutiny. So, no plausible account of knowledge attributions can hold both that speakers are semantically incompetent about “S knows that p” and that the ordinary concept of knowledge is a non-skeptical one, at least not unless the non-skeptic has a radically different view about philosophical analysis. I will argue that skepticism has an error theory that is in many respects at least as plausible as competing accounts of “S knows that p,” and in other respects actually better than competing accounts. Although I accept the analyses of knowledge and justification defended earlier in the dissertation, the arguments I offer for a skeptical interpretation of knowledge attributions apply not only to my favored accounts, but to any similar skeptical position.

In order to make good on this claim, a few things are necessary. I will address them in the following order. I first explain what the problem is. Second, I explain what the competing accounts of “S knows that p” are and what each has to say about the problem of false knowledge attributions. Third, I offer a framework from which to address the problem and assess competing solutions. Fourth, I present the details of the skeptical invariantist error theory I prefer. Fifth, I conclude with a comparative assessment.

#### **4.1 The Problem of False Speech**

Philosophers are attempting to provide an account of the ordinary concept of knowledge. The way we do this is, in part, to think about what ordinary speakers take knowledge to consist in. And one way we do that is to pay close attention to the circumstances under

which speakers are inclined or disinclined to attribute knowledge. The commonsense epistemologist will quickly point out that people say they and others know all sorts of things. The commonsense epistemologist insists that the pervasive tendency to make positive knowledge attributions must not be ignored when providing an account of knowledge. And it is not just a matter of people *talking* a certain way. People really seem to take themselves and others to know certain things. They seem to actually believe such knowledge attributions. Often, when I make an attempt at humor, I am under no pretense whatsoever that what I say is true. Talk about knowledge is usually not that way. At the moment of making a knowledge attribution, speakers are often inclined to take their knowledge attribution to be true. So, the skeptic who takes himself to be providing an analysis of the ordinary concept of knowledge has to explain both why it is that ordinary speakers in possession of the same concept of knowledge as the skeptic make claims that are false by the skeptic's lights, and why it is that those speakers regard their false claims as true. The explanation we seek is what I shall call a "knowledge-attribution skeptical error theory", or simply an "error theory".

We will see shortly that understanding an error theory, and identifying those who need one, is somewhat complicated. So, I delay discussion of that until we have introduced competing accounts of "S knows that p."

#### 4.1.1 *'Knowing' and Knowledge*

A glimpse at much of the contemporary literature in epistemology—perhaps due largely to the popularity of epistemic contextualism—reveals that at some point philosophers started talking a lot about 'accounts of "S knows that p"' and less about 'analyses of

knowledge.’ Philosophers are not always perspicuous about what the relationship is supposed to be between “S knows that p” and S’s knowing that p. Some positions in the debate present themselves as semantic theses—theses about the meaning of sentences of the form “S knows that p.” Contextualists are generally clear about this (Cohen 1999, DeRose 1995; for discussion, see Pritchard 2002). One gets the impression that these philosophers take themselves to be providing the meaning of a particular kind of sentence, not an analysis of knowledge. More classical treatments represent themselves as analyses of concepts or some other non-linguistic entity (Ayer 1956). On the face of it, an account of “S knows that p” and an account of knowledge would never be in conflict.

Connecting these seemingly unrelated issues requires attention to the truth-conditions for knowledge attributions. While one who claims to provide an account of “S knows that p” sounds like he intends only to argue for a view about what propositions are expressed by sentences of that form, at the end of the analysis the philosopher talking about those sentences has to say something about the way the world has to be to make true the propositions expressed by sentences of that form. So while the semantic thesis about knowledge attributions and the metaphysical thesis about knowledge itself are conceptually distinct, one cannot provide a complete account of the semantics of “S knows that p” without eventually getting to the metaphysics of what makes such sentences true or false. It is partly for that reason that Ernest Sosa has argued that the importance of contextualism in epistemology is limited (2000).

So we distinguish (1) knowledge attributions, which are sentences of a certain type; (2) the propositions expressed by knowledge attributions; (3) the way the world has to be to make true or false the propositions expressed by knowledge attributions. So long as there are principles linking knowledge attributions to the propositions they express—so-called disquotational principles<sup>98</sup>—it should be clear that philosophers talking about “S knows that p” are not talking past those talking about the conditions under which S knows that p. Thus, while the positions we consider describe themselves differently—some as explicitly semantic theses, some as metaphysical theses—we see that they are in fact positions in the same debate. What the linking principles are that connect knowledge attributions to propositions expressed by those sentences is a controversial matter. But that there have to be some such principles to make sense of why the semantics of ‘S knows that p’ is relevant to epistemology should be clear.

#### 4.1.2 *Candidate Accounts*

There are several views of the semantics of “S knows that p” on the market today. Since my goal is to show that a skeptical view can compete with all of them, I will describe the competing views in such a way as to highlight their dissimilarities from the skeptical view defended here while largely ignoring their dissimilarities from each other. My

---

<sup>98</sup> See Hawthorne (2004, p.101) for an argument that contextualism cannot accommodate an intuitively clear disquotational principle for ‘knows.’

(Schaffer 2004b, p. 143) claims contextualism does not allow disquotation for ‘true’ because when p is a knowledge attribution the principle “‘p’ is true if and only if p” fails: p may be true in one context but not in another. I think it is fairer to the contextualist to understand the failure of disquotation only for knowledge attributions. So, the expression that lacks a straightforward disquotational principle is ‘S knows that p’; thus, the contextualist’s thorn is the apparently plausible disquotational principle “‘S knows that p’ is true if and only if S knows that p.” In a reply to Hawthorne, DeRose argues that the strength of his contextualist cases makes that candidate disquotational principle less plausible than it first appears (2006).

dialectical interests favor laying out the field with skepticism on one side and everybody else on the other side. Defenders of non-skeptical views have argued at length over the details of similar accounts, but these details are largely irrelevant to my concerns. That is because a skeptical view rejects a core claim shared by most of these accounts, namely, that the truth conditions of “S knows that p” vary across different contexts. Now for the views.

#### 4.1.2.1 *Contextualism*

Epistemic contextualism claims sentences of the form “S knows that p” express different propositions in different contexts of utterance (Cohen 1986, 1988, 1998, DeRose 1992, 1995, Lewis 1996). The truth-value of a knowledge ascription in a context is determined in part by the knowledge ascription’s context. The relevant context is the context of the attributor, not that of the subject. So if S asserts “S knows that p,” S’s first-person knowledge attribution has its truth-value determined in part by S’s context. If J asserts “S knows that p,” the relevant context is J’s.

Contextualism has been defended in part by arguing that the naturalness of knowledge attributions made in some hypothetical conversations is best explained by the hypothesis that the participants in those conversations are speaking truly. Keith DeRose’s bank cases and Stewart Cohen’s airport case both illustrate this (Cohen 1999, DeRose 1992). In Cohen’s airport case, two passengers (Mary and John) are at the airport wondering whether the flight they plan to board has a layover in Chicago. They observe another passenger (Smith) look at his own itinerary and claim to himself that he knows the flight stops in Chicago. Mary and John have an important meeting in Chicago.

Mary says to John, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John conclude that Smith does not know there is a layover in Chicago and decide to answer their own doubts by checking with a gate attendant for more information about the flight.

According to Cohen, Smith is using a lower standard of knowledge than Mary and John because the standards applicable in Smith’s context are lower than those in Mary and John’s context. The intuition that both Smith’s self-attribution of knowledge that the flight stops in Chicago is true and that Mary’s denial that Smith knows the flight stops in Chicago is also true is best explained, Cohen says, by the hypothesis that different propositions are expressed by these knowledge attributions. This version of contextualism would have us understand Smith’s claim, “I know the flight stops in Chicago,” as expressing the proposition *Smith knows<sub>low</sub> the flight stops in Chicago*, and Mary’s claim, “Smith doesn’t know the flight stops in Chicago,” as expressing the proposition *It is not the case that Smith knows<sub>high</sub> the flight stops in Chicago*, where subscripts indicate the knowledge-property expressed by “knows” in the respective contexts of attribution. The operative standard of knowledge in the context in which the knowledge attributions are made partly determines the semantic content of the utterance.

Spelling out exactly what an epistemic standard-determining context is is no small task. We have at least a rough notion of what a context is, though. We attribute ‘tallness’ more generously when we are not discussing NBA players. We are inclined to be more lax with talk about knowledge when we are not answering challenges to knowledge attributions or discussing epistemological issues. And so on. According to

contextualism, the dialectical setting in which we apply these varying standards for the predication of a particular term are distinct contexts: and because contexts set the standard that determines the denotation of a context-sensitive predicate, varying standards of “S is X” or “S X’s that p” yield distinct propositions that are picked out by these expressions in different contexts.

Putting all this together, the contextualist view, in outline, is as follows. Contexts of use set the epistemic standards that are necessary and sufficient for an utterance of “S knows that p” to be true (or false) in that context. The context that determines the truth-value of a particular utterance of “S knows that p” is the context of the knowledge-attributor. So, the proposition expressed by a particular knowledge-attribution is determined by the meaning of “S knows that p” plus the context in which that sentence is uttered.

A few important features of contextualism warrant further comment. First, the claim that it is the *attributor’s* context that determines which proposition a knowledge attribution expresses, and thereby what the epistemic standard is according to which the knowledge attribution will be judged, has led the view to be labeled “attributor contextualism.”

Second, reliance on the attributor’s context rather than the subject’s context or some third-party’s context suggests cases of knowledge attributions made by a speaker regarding a subject not in the context will bring out the distinctive features of attributor contextualism. DeRose (2005) insists on the use of third-party cases even though he

introduced his brand of contextualism using his first-person ‘bank cases.’ (This issue is why I introduced contextualism with Cohen’s third-person airport case.)

Third, contextualists must specify the mechanism responsible for the context-sensitivity of knowledge attributions. For example, many people think ‘flat’ is a gradable adjective: an adjective the proper application of which comes in degrees. Those who hold this view will point out that we readily call a field ‘flat’ even when it has several bumps, so long as none of them would prevent us from playing a competitive game of soccer; but we would not call a table ‘flat’ if it had the same proportion and scale of bumps if we wanted to use it for a desk.<sup>99</sup> On this view, it is a feature of ‘flat’ that it is truly predicated of some surfaces but not others. (Stanley (2005) ch. 2 argues that ‘knows’ is not gradable, perhaps unlike ‘flat.’) It remains an open question, if ‘flat’ is gradable, what explains the gradable character of flatness ascriptions. Perhaps it is because flatness ascriptions include hidden reference to a standard: when I say of the field that it is flat I mean *flat-for-a-soccer-field*, and when I call the table flat I mean *flat-for-a-desk*.<sup>100</sup> Then to predicate ‘flat’ of the soccer field and of the table, even though

---

<sup>99</sup> I side with Unger (1971), (1975) on “flat”: I think “flat” denotes an absolute limit and the contextual variability in conditions under which speakers call something “flat” are best given a pragmatic analysis. As evidence, it is natural after comparing two surfaces for flatness to point to the bumpier one and say, “That one’s not flat. *This* (less bumpy one) is flat.” Repeated application of this practice reveals an inclination to hold that if two surfaces vary in flatness, the less-flat one isn’t flat. Thus, flatness does not come in degrees. Something can be closer to or further from (absolutely) flat, but in a way dissimilar to the way in which something can be more-or-less heavy, or more-or-less hairy.

<sup>100</sup> Note that the clarification of “flat” as “flat-for-an-x” presupposes that “flat” in the clarifying sentence has some meaning on its own, for one could not further clarify “flat-for-an-x” with some further sentence that itself includes “flat.” It follows that even context-sensitive words must have some context-insensitive—that is, invariant—meaning. Without such invariant meaning, the term could not have any cognitive significance for the speaker. This is not itself an objection to contextualism. The contextualist could happily agree that *epistemic* terms possess invariant meaning but that the use of such expressions in different contexts yields different propositions. While that reply is available, some feature of language still

they clearly vary in how flat they are, is to compare the surfaces according to different standards. There is debate about how to understand gradability as well (cf. DeRose (2008)). The issues that have to be addressed to achieve a fully fleshed-out contextualism about 'flat' are many and complicated; analogous issues are pressing for the epistemic contextualist.

Fourth and finally, none of this addresses the metaphysical question of what constitutes knowing a proposition in a context. Contextualists have different accounts of this. (Dretske 1981) and Lewis (1996) have both defended a relevant alternatives account according to which knowing that p requires eliminating all relevant not-p possibilities while context determines which possibilities are relevant. Cohen (1988) holds that context determines the standard of justification necessary for knowledge. DeRose (1995) claims context determines how strong the sensitivity relation has to be in order to know. Clearly, contextualism allows for variation on the issue of what constitutes knowing in a context.

Fortunately, in order to address the comparative merit of epistemic contextualism as such versus the other views it is not necessary to consider every possible version of contextualism. Only the fundamentals are necessary for a comparative assessment.

#### 4.1.2.2 *Invariantism*

Because many of the terms and distinctions needed to introduce the views relevant to this debate were introduced by way of presenting contextualism, the following introductions

---

has to be responsible for context-sensitivity. If it is not a feature of the epistemic terms themselves, then what could it be?

can be briefer. Invariantism, in all its forms, is the denial of contextualism. Thus, according to invariantism, the proposition expressed by a knowledge attribution is fixed once ‘S’ and ‘p’ are fixed. A knowledge attribution expresses the same relation to a proposition in every context of use. Different versions of invariantism emerge when accounts of “S knows that p” are offered. The two species of invariantism are strict invariantism and impure invariantism (Fantl and McGrath 2009, Gerken 2011, Hawthorne 2004, Stanley 2005).<sup>101</sup> Strict invariantism holds that the truth-value of a knowledge attribution depends only upon traditional epistemic factors like the subject believing the relevant proposition and having an adequate basis for belief in the proposition. Impure invariantism denies that only traditional epistemic factors determine the truth-value of a knowledge attribution. Impure invariantists allow, most importantly, that how important it is to S that p is true can impact whether or not S knows that p while other factors are held fixed, such as p’s truth, S’s belief that p, and S’s justification for believing that p.

Moderate invariantists are those who hold a semantics for “S knows that p” which results in ordinary knowledge attributions often being true (Hawthorne, 2004).<sup>102</sup>

Skeptical invariantists hold that “S knows that p” is very often false, perhaps even mostly or always. Since ‘moderate’ and ‘skeptical’ are labels describing how many knowledge

---

<sup>101</sup> “Impure” is the name given the view by Fantl and McGrath. They defend the view. Other names are “Interest-Relative Invariantism” (Stanley) and “Subject-Sensitive Invariantism” (Hawthorne).

<sup>102</sup> In this way of describing the view, what makes moderate invariantism ‘moderate’ is not its epistemic standards, but that it implies that many knowledge attributions are true. (Hawthorne (2004) characterizes the position this way.) It is consistent with a view being ‘moderate’ invariantist that it accepts very high standards for knowledge, provided the standards are met regularly.

attributions come out true on a given semantics for “S knows that p,” one could consistently be a moderate impure invariantist or a moderate strict invariantist; similarly, one could consistently be a skeptical impure invariantist or a skeptical strict invariantist.

#### 4.1.2.3 *Contrastivism and Relativism*

According to Mark Richard, relativism is implied by contextualism (2004).<sup>103</sup> Richard’s main argument seems to be that one cannot make sense of how surface-contradictory ascriptions of “X is F” can both be true unless the *property* of being an F is a relative property (cf. p. 232). For knowledge attributions, then, “S knows that p” and “S does not know that p” can both be true only if *knowing that p*—that is, *knowledge*—is relative. Understood this way, epistemic relativism is disquoted epistemic contextualism. The world has to be a certain way for knowledge attributions to have any truth-value at all. Relativism says the world has to have relative properties for surface-contradictory knowledge attributions to both be true.

Like Richard’s relativism, contrastivism picks up where generic contextualism leaves off. In the words of Jonathan Schaffer, contrastivism’s main defender, “Contrastivism treats ‘knows’ as denoting a ternary relation with a slot for a contrast proposition” (2004a, p. 73). In this way, contrastivism provides a specific account of why knowledge attributions are context-sensitive. Contrastivism holds that “S knows that p” makes hidden reference to a contrast proposition q. So the truth-value of “S

---

<sup>103</sup> Richard is not alone in defending relativism, though he is perhaps the earliest contemporary defender of the view that pro-contextualist arguments actually support relativism. See also MacFarlane (2007), (2009). For a contrary view, see Davis (2013).

knows that p” depends in part upon whether S knows that p *rather than q*. As an example, Schaffer points to the disagreement between the skeptic and Moore as to whether Moore knows he has hands. A way to allow that the skeptic truly asserts “Moore does not know he has hands” while Moore truly asserts “I know I have hands” is to point to different contrast propositions: the skeptic is right that Moore does not know he has hands *rather than that he is a handless brain-in-a-vat*, and Moore is right that he knows he has hands *rather than elbow-length arm stumps* (2004, p. 80). Now if the attributor’s context determines which contrast proposition is implicit in a particular utterance of “S knows that p,” then contrastivism will deliver identical verdicts on cases contextualism does: contrastivism will just be a particular way of spelling out contextualism because it specifies a partial mechanism for explaining context-sensitivity. In fact, though, Schaffer has little to say about how contrast propositions are selected when they are not selected explicitly, e.g. by someone asserting “S knows that p *rather than q*.”

I introduce all of these views for the sake of completeness. In much of what follows I will say little about contrastivism, relativism, or impure invariantism. It should be clear where insights those views offer impact what can or has to be said about the issues I focus on. But for the sake of my main argument, I will focus on strict moderate invariantism and contextualism on one side and skeptical invariantism on the other. This choice does not treat these other views unfairly, for if the defenders of contrastivism and relativism are right, whatever can be said for or against contextualism usually goes for the other views as well. For the same reason, I will largely ignore impure moderate invariantism (Hawthorne 2004, p. 157). So when I say, “the truth-maximization principle

of charity favors contextualism here,” relativists, contrastivists, and impure invariantists, by their lights, accrue that benefit as well. They accrue the benefit because all of those views allow that surface-contradictory knowledge attributions can both be true in some circumstances. The choice to limit the theories under consideration is driven by my interest in exploring the support that can be claimed for a non-skeptical interpretation of knowledge attributions over a skeptical one. So for ease of exposition, I will focus mostly on contextualism and strict moderate invariantism in comparison with skeptical invariantism.

#### 4.1.3 *Semantic Blindness and Linguistic Sensitivity*

Some of the most powerful objections to contextualism in the literature concern the ignorance of language contextualism requires of speakers. Let us call this kind of error semantic blindness (following Hawthorne (2004)). To see this, consider the so-called skeptical ‘paradox.’ According to Cohen’s (1988) resolution of the skeptical paradox, all three of the following sentences can be true:

1. “I know I have hands.”
2. “I do not know I am not a handless brain in a vat.”
3. “I know I have hands only if I know I am not a handless brain in a vat.”<sup>104</sup>

According to Cohen, we worry about the skeptical threat posed to (1) by (2) plus (3) only because we fail to appreciate the fact that these claims, when true, are true *in different*

---

<sup>104</sup> Though it is not always made explicit in the literature, contextualists need to present the problem as an issue between apparently inconsistent *sentences* rather than propositions, since on the contextualist view the propositions are consistent.

*contexts*. So, (2) plus (3) being true in high-standards contexts is no threat to (1) being true in low-standards contexts. While contextualism successfully captures the plausibility of each of these claims considered individually, it carries a cost. Contextualism can offer its solution to the skeptical problem only by maintaining that speakers systematically confuse the propositions expressed by knowledge attributions made in some contexts with propositions expressed by knowledge attributions made in other contexts (Schiffer 1996). Once a speaker is in a high-standards context, any previous assertion of (1) seems false, and “I do not know I have hands” seems very strongly to be a denial of (1). But contextualism denies that retracting (1) in that way is possible: if (1) was asserted under low standards, the assertion of “I do not know I have hands” under high standards will not deny (1), but will say something irrelevant to (1).<sup>105</sup> In that way, the contextualist resolution of the skeptical paradox depends upon speakers being insensitive enough to contextual shifts affecting sentence meaning such that they fail to realize that (1), (2), and (3) are compatible, while being sensitive enough to contextual shifts affecting meaning to say “I don’t know that p” upon being pressed.

While there are some cases where surface-contradictory sentences turn out to be compatible, speakers made aware of the propositions expressed by those sentences are usually made aware of the sentences’ compatibility easily. Consider: I might call a friend ‘tall’ when he can reach something on top of the refrigerator that I cannot, but deny that he is ‘tall’ during a conversation about the upcoming NBA draft. Yet no one present for

---

<sup>105</sup> Of course, contextualism does not require that “I do not know I have hands” is *never* a denial of (1): the two are incompatible when they are asserted in contexts with shared epistemic standards.

both conversations would take my refusal to attribute tallness in the draft context to contradict my earlier attribution of tallness in the kitchen—and if they did, any appearance of contradiction disappears when I clarify that by ‘tall’ in the kitchen context I meant ‘tall-for-a-normal-person’ and that by ‘tall’ in the draft context I meant ‘tall-for-an-NBA-player.’ But when one speaker says “S knows that p” and another says, of the same S and p, “S does not know that p,” we hear these as incompatible and the apparent contradiction is not easily dissolved by pointing to context-sensitive standards: “I meant ‘knows’ by *normal* standards...”<sup>106</sup> Since apparently contradictory knowledge attributions are not easily shown to be compatible by reference to distinct standards of evaluation, unlike statements involving less controversially context-sensitive terms like ‘tall,’ contextualism implies that speakers are blind to the context-sensitivity of ‘knows’ (Cf. Rysiew 2001, p. 484).

So if contextualism is true, it is a surprising fact that speakers demonstrate context-insensitivity in this unique way. It should be clear that this objection extends to other views introduced in the last section. Contrastivism allows for context-variability in a way parallel to the context-variability allowed by contextualism because contrastivism provides a specific mechanism for context-variability. Thus the objection from context-insensitivity extends to contrastivism as well. Speakers hear “S knows that p” and “S does not know that p” as contradictory, even if the former sentence expresses *S knows*

---

<sup>106</sup> But see Ludlow (2005) for discussion of alleged standards-markers. Note that I haven’t claimed there are no standards-markers, but that the apparent incompatibility of knowledge ascriptions is not easily dissolved by their use in a way similar to how the apparent contradiction between height ascriptions is easily dissolved by saying, “I meant tall-for-a-[insert relevant standard].”

*that p rather than q* and the latter sentence expresses *S knows that r rather than p*.

Similarly, because epistemic relativism is disquoted contextualism, the objection from context-insensitivity applies there, too.

According to impure invariantism, factors such as how much it matters for one that one's belief is true partly determines whether or not one knows. Consider again the airport case. Smith, apparently in no rush, claims to know the flight stops in Chicago. Mary and John, anxious to get to Chicago for an important meeting, claim not to know. Let us suppose that Mary, John, and Smith all have the same epistemic standing with respect to the proposition *the flight stops in Chicago*. Smith, Mary, and John all have what would normally pass for solid evidence: a printed itinerary that says the flight stops in Chicago. Impure invariantism can allow that Smith knows the flight stops in Chicago while Mary and John do not know, while the only relevant difference is how important getting to Chicago is for the parties, respectively.

Where does semantic blindness come in? If impure invariantism is true, it is possible to truly say things like this: "I know the flight stops in Chicago, but if I'd been worried about it, I wouldn't have known," and "If only I wasn't so worried about getting to Chicago, I'd know I'd get there." Those claims sound false, regardless of who one imagines saying it. That suggests there is something very strange about the view that allows for those claims to be true. That is not a knock-down argument to the view. But semantic blindness objections pose serious problems for the philosopher claiming that the correct way of understanding ordinary knowledge talk is by appealing to implicit context-sensitivity, for claims like the above sound far from acceptable.

While the problems facing contextualism are well-known, what seems not to have been noticed is that strict moderate invariantism—probably the majority view among epistemologists—faces difficulty explaining why speakers retract or conditionalize true knowledge claims in contexts where error possibilities have been raised. (To conditionalize a claim is to repeat what one has already said as the consequent of a conditional statement with some new antecedent specified. For example, I might say, “It’s going to rain today.” A friend challenges me: “You know weather forecasts are often unreliable, right?” I reply, “I mean *if* the forecast I read is accurate, *then* it’s going to rain today.”) According to strict moderate invariantism, speakers in such contexts often know the proposition in question. Yet when some error possibility is brought to their attention, the obtaining of which would *not* falsify the original claim, they often retract or conditionalize the knowledge claim. If strict moderate invariantism is true, it may therefore seem that speakers are insensitive to what is and is not *relevant* to the truth of a knowledge attribution. Call this linguistic error *relevance blindness*.

To see this, consider again the airport case. Most strict moderate invariantists will hold that Mary and John *do* know the flight stops in Chicago.<sup>107</sup> Their worry about their flight arriving on time and in the right place does nothing to undermine the traditional epistemic factors that support their belief that the flight stops in Chicago: factors like belief, truth, and adequate justification. On this view, the misprint hypothesis—being the

---

<sup>107</sup> An exception might be Bach, who might be read as claiming that Mary and John’s worries about the flight cause them to lose the belief that the flight will stop in Chicago, and consequently not to know it (2005). I won’t engage with Bach here except to mention that the most natural reading of the airport case is that the parties believe the proposition but are battling off worries that prompt them to claim they and others relying on the itinerary don’t know.

sort of ‘far fetched’ scenario that would only unfold in a distant possible world—is not *relevant* to whether or not John and Mary know the flight stops in Chicago. But if that is correct, moderate invariantists have to explain why Mary and John *believe* the misprint scenario is not only relevant to whether or not anyone reading the itinerary knows the flight stops in Chicago, but so relevant that they think no one *can* gain knowledge of the flight’s destination by looking at their itinerary. Generalizing, if moderate invariantism is true, the pervasive tendency to retract, conditionalize, or even defend knowledge claims against any-and-every error possibility that becomes salient must be due to speakers’ ignorance of what is and is not relevant to the truth of knowledge attributions.

When speakers are competent users of some term, they are generally sensitive to what is and is not relevant to the correct application of that term. The sensitivity in question need not be thought of in cognitive terms. We don’t have to think of speakers as knowing or believing *that* X is relevant to Y in order for speakers to be sensitive to the relevance of X to Y. The kind of sensitivity in question can be understood solely as a disposition to speak and behave in certain regular ways. A competent user of ‘tiger’ is sensitive to core features relevant to whether or not ‘tiger’ is applied correctly (Putnam 1975, p. 250). If S is a competent user of ‘tiger’ and asserts of a nearby object, ‘X is a tiger,’ S will *not* retract that assertion if someone else asserts, ‘There are ants nearby.’ That is because the presence of ants is semantically irrelevant to whether or not ‘X is a tiger’ is true, and those competent with ‘tiger’ are sensitive to that semantic fact.

Hawthorne introduced ‘semantic blindness’ as a label for the sort of ignorance of language of which context-insensitivity is a kind (2004).<sup>108</sup> Relevance blindness is plausibly another kind of semantic blindness. It is unclear whether we should think of relevance blindness as more, less, or comparable in severity of error to context-insensitivity. A justification for taking semantic blindness objections seriously—either as a problem for an opponent or a serious enough problem for one’s own view that it deserves a response—is rarely provided. But it seems that something like the following is assumed by those who offer and respond to semantic blindness objections. A *methodological principle of charity suggests that we ought to regard speakers as generally sensitive to the semantics and pragmatics of familiar terms in their own languages.*<sup>109</sup> The idea should be clear: competent speakers using familiar terms are generally sensitive to the proper behavior of those terms. Those competent with ‘tiger’ are generally sensitive to the proper behavior of ‘tiger.’ They will recognize as infelicitous “I ran tigerly today” because they know ‘tiger’ is not used adverbially; “Tiger rules the jungle” because they know ‘tiger’ isn’t a proper name; and will not take “There

---

<sup>108</sup> The issue of “semantic blindness” seems to me to cover a host of issues regarding how to correctly talk about knowledge - issues which, if contextualism is true, speakers are severely mistaken about. Too often, the different forms of ignorance required by contextualism are not kept apart, which might have something to do with why some philosophers think semantic blindness is so damning to contextualism (e.g. Schiffer 1996) and some think it isn’t problematic at all (DeRose 2006). My discussion specifically concerns semantic blindness and relevance blindness as species of linguistic insensitivity.

Abath (2012) has argued, that whether or not semantic blindness is a problem for contextualism, what Abath calls “content unawareness” is. In short, if contextualism is true, people don’t know what propositions their sentences express. On my way of categorizing things, his content unawareness is a kind of semantic blindness.

<sup>109</sup> This should not be taken to suggest that speakers easily *distinguish* the semantic content of an utterance from its pragmatic content without some guidance. One can be sensitive to A and B without being sensitive to whether something is an A rather than a B (or vice versa). E.g. while sitting in a chair one can be aware of pushes and pulls in the chair without being aware of a given force’s being a push rather than a pull (or vice versa).

are ants nearby” to challenge the assertion (made while pointing) “*That’s a tiger.*”<sup>110</sup>

Examples abound.

We generally take for granted that speakers know how to use familiar terms in their own languages. It is tempting to identify speaker linguistic sensitivity to the behavior of a term with semantic competence with that term. Someone failing to recognize the irrelevance of “There are ants nearby” to “*That’s a tiger*” suggests that the speaker *does not know* what ‘tiger’ *means*. Similarly, suppose someone thought “I am here” always contradicted “I am not here,” regardless of who is speaking and where: so in this person’s mind “I am here” uttered outside of the library literally contradicted “I am not here” uttered by someone else, somewhere else. We would be inclined to say this person just does not know what ‘I’ or ‘here’ *means*. The presumption of linguistic sensitivity suggests that if a speaker is sensitive to the correct behavior of ‘tiger’ and ‘I’ or ‘here,’ the speaker will not make the kinds of mistakes just described. So the principle would have us conclude that these speakers are not competent with those terms.

In sum, let us say the following about linguistic sensitivity.

(1) We expect that speakers competent with a term T are sensitive to the behavior of T. That includes:

- (1a) sensitivity to the truth and falsity of expressions using T when in possession of adequate background information pertaining to whether or not something is a T,
- (1b) sensitivity to the appropriateness of expressions using T,
- (1c) acceptance of core conventionally-determined T-relevant

---

<sup>110</sup> None of these examples are instances of *necessarily* infelicitous speech that would be recognized as such by competent speakers. One’s linguistic community, even a sub-community such as a group of intimates, *could* accept a convention that allowed for a metaphorical use of ‘tiger’ as an adverb, so that “X-ed tigerly” means “X-ed in a powerful/stealthy/some other way”; and a community could agree to the convention that use of ‘tiger’ as a proper name is a personification of the class of tigers. The fact that the above uses are infelicitous is thus contingent. And so it is for all uses of all expressions, because it is in every case a contingent fact that speakers adopted the convention to use ‘t’ for t.

features (e.g. those competent with ‘tiger’ accept that typical tigers are striped).

(2) We do not expect that speakers competent with T will *never* make mistakes regarding (1a-c). But mistakes will be the exception rather than the rule, and will not include *systematic* mistakes concerning *core* T-relevant features (cf. (1c)).

(3) We do not expect that speakers competent with T are capable of providing a definition of T.

Regarding (2), it is not possible to define in the abstract what ‘core’ features a speaker must be sensitive to in order to count as being competent with a given term in a language. What makes a feature a core feature is just that we count speakers as being competent (or incompetent) with a certain term if and only if they accept (or fail to accept) certain claims made using that term. We generally agree about how to decide individual cases. E.g. we all agree that a speaker who accepts “Tigers are *not* striped” is not competent with ‘tiger’ (cf. Putnam 1975, p. 250.)<sup>111</sup> But it is overwhelmingly obvious that there are core T-relevant features to which speakers have to be sensitive if they are to count as T-competent.

The presumption of linguistic sensitivity should seem like a familiar notion, and one we ordinarily assume is true. The presumption, simply, is that competent speakers

---

<sup>111</sup> The presumption of linguistic sensitivity has been objected to on the basis that semantic externalism is consistent with speakers being mistaken about what sorts of things are Fs while they remain competent users of ‘F.’ I’ve used Putnam’s tiger example throughout for a reason. The arch-externalist Putnam required for semantic competence that speakers be adept with a term even if they sometimes made mistakes with the term. If the requirement is compatible with Putnam’s externalism, then it should be for other versions of externalism, too. What prevents a speaker from being competent with a term the speaker regularly uses is the making of *systematic* mistakes regarding *core* T-features, which is just what we required above in condition 2. One need not be semantically omniscient or infallible to satisfy my conditions for semantic competence. (I should note that I’ve received the objection that other externalists, notably Burge, have presented stronger cases of semantic insensitivity: e.g. a man who incorrectly claims to have arthritis in his thigh (1979). The objection is, I believe, irrelevant, for there is no plausibility *at all* to the idea that someone who claims to have arthritis anywhere other than a joint is competent with “arthritis.”)

know how to speak competently and that when they do make a mistake, either the mistake is not major or it is easily correctable.

The hypothesis of speaker context-insensitivity of ‘knows’ violates the presumption of linguistic sensitivity because speakers supposedly competent with ‘knows’ actively resist the contextualist explanation of the skeptical paradox that “S knows that p” does not always contradict “S does not know that p.”<sup>112</sup> Relevance blindness for familiar terms, including ‘knows,’ would also violate the presumption, because speakers supposedly competent with ‘knows’ would fail to be sensitive to what is and is not relevant to the truth of a knowledge attribution. When a philosophical thesis violates the presumption of linguistic sensitivity, some explanation is needed: either an explanation of why, contrary to appearances, the account in question is consistent with the presumption, or of why the proposal in question is a justified exception to the presumption.

#### **4.2 Methodological Assumptions for the Analysis of “S knows that P”**

An objector might simply concede the point that the views described above all require some ignorance of language that could be called linguistic insensitivity. But so what—the objection goes—we all learned from Davidson (1973) that when you want to figure out what an expression in a language means you presuppose that the expression is often used to express a true proposition. Skepticism is fundamentally at odds with this presumption because even a modest skepticism entails “S knows that p” nearly always

---

<sup>112</sup> But see DeRose (2006) for an argument that contextualists and invariantists share the problem of semantic blindness.

expresses a false proposition. So skepticism is at odds with the interpretive principle respected by non-skeptical views about the semantics of “S knows that p.” Some such interpretive principle is necessary. To figure out what “S knows that p” means we must act as interpreters of our own language, and to do that we must have some methodological assumption that will aid us in figuring out what “S knows that p” means. A principle of charity looks like the perfect candidate. But it threatens—the objection continues—to rule out skepticism as a non-starter.

I take this objection seriously. I am sympathetic to the idea that underlies it, namely, that epistemologists should be trying to figure out what the correct analyses of concepts of epistemological interest are, and that we do that at least in part by careful attention to language. Consequently, the epistemologist who wants to interpret sentences of the form “S knows that p” in order to arrive at an analysis of knowledge needs to be equipped with some methodological assumptions that determine what analysis of knowledge can be inferred from certain patterns of use. In this section, I explore one such principle—a principle of charity—and consider how it can help resolve this debate. The argument of this section does not allow for a decisive victory of skepticism over its competitors. However, I believe I can show that considerations of charity do not favor non-skepticism over skepticism: indeed, I will argue on the basis of considerations of linguistic sensitivity that a skeptical interpretation is supported by ordinary usage of “S knows that p.”

#### 4.2.1 *The Truth-Maximization Principle of Charity*

It is assumed that being charitable in one's interpretation of another's words is a good thing. No one would accuse an opponent of offering a charitable interpretation with the intention of maligning the opponent's work. But what could be good about a theory that is charitable? The natural answer is that a charitable interpretation is, *ceteris paribus*, more likely to be *true* than an uncharitable interpretation. So understood, a principle of charity is a principle of theory selection. When considering alternative theories about the semantics of an expression, the charitable theory is preferable to its competitors because it is more likely to be true than its competitors. So one ought to be charitable because charity is truth-conducive. It is hard to imagine what other motivation there could be. Surely philosophers do not identify a thesis as charitable simply to say something nice about the thesis independently of an appraisal of the likely truth of that theory. So we will assume that a principle of charity is a principle of theory selection because charitable theories are, *ceteris paribus*, more likely to be true.

Donald Davidson's principle of charity was invoked as a methodological constraint on the interpretation of other speakers' utterances (1973). Davidson thought that figuring out what a speaker means by an expression requires assuming that much of what the speaker says using the expression is true. By doing this, Davidson thought the interpreter would be in the best position to correctly attribute beliefs to the speaker. Davidson's motivation for the principle of charity was his project of radical interpretation, in which an interpreter attempts to understand a language without any previous knowledge of the language. Semanticists of epistemology obviously know the

language they are working with; but while they do not have precisely the same motivation behind their *use* of a principle of charity as Davidson did, the version of a principle of charity many rely on is, for all important purposes, the same. We can state that version as follows: when constructing a semantics for ‘knows’ (and its cognates, and perhaps other important epistemic terms like ‘rational’ or ‘justified,’ for which an analogous principle is easily constructed), one ought *ceteris paribus* to offer a semantics that makes a very high percentage of knowledge attributions come out true.

Evidence for the reliance on this version of a principle of charity is largely implicit, though it has been explicitly called upon as a motivation for some views. In the quotes below, DeRose and Montminy, both contextualists, claim charity favors their view. Brown, a strict moderate invariantist, claims charity favors contextualism, and Williamson, another strict moderate invariantist, gestures toward considerations of charity favoring contextualism and impure (‘sensitive’) invariantism.

But I think the reason that helps in supporting the claim that what one’s imagined speaker is saying is true is that it engages the general presumption that where speakers are not basing their claims on some false beliefs they have about underlying matters of fact, *how they naturally and appropriately describe a situation, especially by means of very common words, will be a true description.*<sup>113</sup>  
DeRose, (2005)

---

<sup>113</sup> DeRose includes as a condition on a principle of charity that the knowledge attributions whose truth is to be assumed not be based on false beliefs about “underlying matters of fact.” While some sort of qualification like this is indeed necessary, it’s tricky to figure out how to state the qualification. However it is put, the principle should be understood as not granting the truth of knowledge attributions which are based on inference from a false premise (as in the stopped clock Gettier case), but how to precisely state the qualification while avoiding triviality remains a challenge: one false belief a subject has might be *that they know the proposition*, but surely the DeRose qualification shouldn’t say, “presume knowledge attributions are true, except when the subject believes he knows that p but his belief is false because he does not know that p.”

One major advantage of [contextualism] is that it respects the context-sensitivity of our *intuitions about the truth values* of such claims. The (moderate or radical) sceptic's knowledge denial is plausible, since it *seems correct to deny knowledge* to a subject who cannot eliminate possibilities of error; yet, in everyday contexts, *we readily attribute knowledge* to subjects who have acquired their beliefs on the basis of fallible evidence.

Montminy, (2009)

Contextualists offer a *charitable* understanding of intuitions [about the correctness of knowledge attributions], interpreting them *as reflecting the truth value* of the knowledge attributions and the appropriateness of the relevant assertions and reasoning.

Brown, (2005)

The common basis of contextualism and sensitive invariantism comprises pairs of claims...which in effect endorse the dispositions of speakers in various contexts to assert or deny 'know' of various cases. Presumably, the endorsement rests on a methodological principle of charity, by which, very roughly, we should prefer to *interpret speakers as speaking and thinking truly* rather than falsely (*ceteris paribus*).

Williamson, (2005)

Philosophers identifying being charitable with preserving the intuitive truth-value of a high percentage of both positive and negative knowledge attributions usually do not state their assumption this plainly. For the sake of argument, I will suppose considerations of interpretive charity are a strong motivating factor behind defenses of contextualism and strict moderate invariantism, rather than the only one, or the most important one, or some such stronger claim.

Contextualism probably fares the best with respect to the principle of charity if charity is understood as truth-maximization of knowledge attributions. (Similar remarks apply to impure invariantism, contrastivism, and relativism.) It is easy to see why.

Recall the skeptical paradox:

1. "I know I have hands."
2. "I do not know I am not a handless brain in a vat."

3. "I know I have hands only if I know I am not a handless brain in a vat."

Contextualism allows that all three of the following sentences are true, just not at the same time. Contextual shifts make (1) and (2) true in some contexts and false in others. Individually, each of (1) and (2), combined with (3), entails the falsity of the other. Speakers can say things like what the skeptical paradox says. Contextualism can make all three sentences true. This is a mark in favor of contextualism by the lights of a truth-maximization principle of charity.

Contextualism is also in the nice position of not needing to maximize the truth of positive knowledge attributions over the truth of negative knowledge attributions, or vice versa. Contextualism can allow that surface-contradictory positive and negative knowledge ascriptions are *both* true most of the time. That is because positive and negative knowledge ascriptions are often compatible with each other after accounting for shifting standards for knowing.

Strict moderate invariantism fares worse than contextualism on this score. Recall the strict moderate invariantist account of about the airport case. There, Mary and John anxiously deny that (calm) Smith knows the flight stops in Chicago while all three of them know the flight stops in Chicago. Mary and John make a false knowledge attribution because they mistakenly think the possibility of a misprint on their itinerary has to be eliminated in order for them to know the flight stops in Chicago. These moderate invariantists hold that a positive knowledge attribution does not require for its

truth that the subject of the attribution be in a position to eliminate all skeptical threats.<sup>114</sup>

But speakers who satisfy moderate invariantist conditions for knowledge often deny knowing when doubts are raised, even though the elimination of the error possibilities raised in the doubts is not necessary for knowledge. The speaker in this position who denies knowing that p thus speaks falsely.

The truth-maximization version of the principle of charity would regard that version of moderate invariantism as less charitable than contextualism because while contextualism allows that both positive and negative knowledge attributions are very often true, this kind of moderate invariantism holds that negative knowledge ascriptions are often false: speakers deny knowledge because they are, in a way, confused.<sup>115</sup> This version of moderate invariantism accepts that a smaller percentage of knowledge attributions are true than contextualism does, and is consequently less charitable by the truth-maximization principle.

Of the views considered, skeptical invariantism fares worst with respect to the truth-maximization principle of charity. Skeptical invariantism derives its name from what it implies, namely, that ordinary knowledge ascriptions are generally false. So it is necessarily the case that skeptical invariantism will be uncharitable if charity is understood as truth-maximization: if charity requires making many knowledge ascriptions come out true, and skeptical invariantism makes them generally false, then skeptical invariantism is uncharitable.

---

<sup>114</sup> Cf. Rysiew (2001), Williamson (2005), Bach (2005), Black (2005), Brown (2006).

<sup>115</sup> Rysiew basically accepts the conclusion that knowledge-deniers in the sort of situation described above are confused (2001). The explanation of why they are confused is the focus of Rysiew's paper.

The truth-maximization principle of charity claims that speakers should be interpreted in such a way as to make a high percentage of their claims come out true. It is unclear just how high is high enough to be charitable full-stop, but it should be clear that if one proposed semantics for a term in a language yields a higher ratio of true claims to false ones than an alternative semantics for the same term in the same language, then the former semantics is more charitable with respect to truth-maximization of utterances using that term in that language. That is what it is for a proposed semantics for a term in a language to be more charitable than another proposed semantics. In the case of a semantics for epistemic terms in a language, the most charitable position is the one that delivers the highest percentage of true claims made using a certain term relative to all other competitors.

The scoreboard at the moment puts on top contextualism and the other views that allow for the contextual variability of the truth-values of knowledge attributions, followed by strict moderate invariantism, with skepticism necessarily far behind the others.

However, there is a problem with the truth-maximization principle of charity. The most serious concern with the principle is its easy misapplication. The principle says *ceteris paribus*, one's semantic theory should be constructed so as to make a high percentage of knowledge attributions true. There are some defeating conditions. The use of exaggeration offers an instance. In English, use of the expression 'a million' is more often than not exaggeration: "The drive here was a million miles," "I dropped a box of toothpicks and had to pick them all up. I swear there were a million of them!" and so on.

If it really is the case that most uses of ‘a million’ are for the purpose of exaggeration, then applying the truth-maximization principle of charity to assertions involving that expression will lead us astray. The principle would have us presume that the correct semantics of ‘a million’ is the one that makes most assertions using ‘a million’ true. Were we do to so, say, by defending contextualism about ‘a million’, then we would understand the expression to *literally* mean, in some ordinary contexts, “an indeterminately or annoyingly large quantity,” and in more precise contexts, “the number identical to 1,000 x 1,000.” That is clearly unacceptable. ‘A million’ means *1,000,000* in any use, which is what makes exaggerated speech *non-literal* speech.<sup>116</sup>

Loose talk offers another counter-instance. Unger (1975) argued that we often make knowledge claims in a loose way. We make claims of the form “X is F” when it is merely the case that X is close enough to being F for present purposes. The assertion is therefore false, but not problematically so, because it does not interfere with our practical or communicative purposes. Davis (2006) offers a nice example of loose talk:

When the scoop comes up empty in the coffee jar, I yell to my wife, “The coffee is all gone.” When my son comes down for breakfast a few minutes later, he announces that he needs a few coffee grounds for his science project, and then asks, “Is the coffee really all gone?” I say with no embarrassment, “No, there may be enough for you.”

---

<sup>116</sup> (Butchvarov 1970) and (Schaffer 2004b) both try to understand false knowledge attributions as exaggerations and exaggerations alone, without calling on any of the additional resources discussed in this section. I think this is a mistake unless “exaggeration” is construed broadly enough that loose uses of language include exaggeration. It seems to me that there is a difference between what is going on when I say (exaggerating), “I had to wait in line behind a million people at the DMV to renew my driver’s license” versus when I say (loosely speaking), “I was going to have dinner with my wife when I got back from the DMV, but the house was empty.” But as my focus is on using language non-literally to make false statements, I agree for the most part with Butchvarov’s and Schaffer’s treatment of the phenomena. I discuss their views further below.

The speaker was trying to express that there was not enough coffee to make a pot of coffee. What he literally and falsely asserted, however, was that there was no coffee *at all*. The claim that the coffee was completely gone was loose speech. Applying the truth-maximizing principle of charity to loose uses of terms would result in semantic trouble. No one should want to defend the semantic theory that says “there is no coffee” is *literally* true when there is still some coffee.

Loose use of language and exaggeration are two uses of language where the truth-maximization principle of charity would suggest that we accept that an expression means that *p* while competent speakers would universally reject that the expression means that *p*, because competent speakers hold that the expressions in question are literally false even while they convey something meaningful. In general, that a principle of theory selection includes a *ceteris paribus* clause is surely no reason to abandon the principle. The same holds here: the truth-maximization principle of charity has known defeating conditions, but that is insufficient to discard the principle. But the presence of such conditions does mean the principle must be applied with care. When it is possible that an expression is being used in a way that speakers of the language would recognize as false speech, independent arguments are needed to show that common uses of a particular expression are not instances of deliberate false speech. When such uncertainties are present one cannot lean on the truth-maximizing principle of charity alone to select the true semantics.

This is particularly important in the case of knowledge attributions because skeptical epistemologists have argued that positive knowledge attributions are false and

recognized as such by ordinary speakers—or at least, speakers are easily capable of becoming aware that positive knowledge attributions are always or nearly always false. According to these skeptics, positive knowledge attributions are often exaggerations or loose ways of speaking (BonJour 2010, Butchvarov 1970, Unger 1975).<sup>117</sup> So the fact that a non-skeptical semantics makes more knowledge attributions true is therefore not decisively a mark in its favor: it may be that knowledge attributions are the sort of expression that are regularly used to express false propositions. As interpreters trying to figure out what “S knows that p” means, we have to take seriously the possibility that knowledge attributions are often false even while they are used to successfully convey thoughts and information to others.

The need to eliminate such possibilities means semanticists who succeed in showing that a proposed semantics for an expression is more charitable in the truth-maximization sense have not thereby shown that the more charitable semantics is more likely to be correct than the less charitable one. A principle of charity is a methodological tool for theory selection. In this context, it helps the interpreter figure out how to approach a given expression in order to figure out what that expression means. The truth-maximization principle of charity says to take for granted that uses of the expression are true, all things being equal. But sometimes things are not equal. The truth-maximization principle of charity is useful for assigning initial plausibility to opposing theories only in the absence of information about how a particular expression

---

<sup>117</sup> Some have discussed the availability of an exaggeration or loose talk error theory without endorsing it (Fumerton 2006b, Hawthorne 2004). Davis defends a loose talk theory in favor of some deliberate false knowledge attributions, but it is less clear that he can be considered a *skeptical* invariantist (2006).

might be used to make false claims. Once we reach the stage where we are considering data regarding how the expression in question is used to make false claims, the work of the principle of charity is done. Since the interpreter has to take into account how the expression is used in concrete scenarios—not only in the absence of information about use, which is where the principle of charity is most helpful—data concerning false uses is vitally important for deciding what the expression means.

Non-skeptical proposals for the semantics of “S knows that p” undoubtedly make more knowledge attributions true than skeptical invariantism, and therefore they are more charitable, in the truth-maximizing sense of ‘charitable’, than skeptical invariantism. But whether the former are therefore more likely to be true than the latter depends on whether non-skeptical semanticists can successfully refute the skeptical semanticists’ arguments that knowledge attributions are often a kind of deliberately false speech and on how the different views fare with respect to the second version of the principle of charity, which I will now introduce. I will argue that skeptical invariantism has the upper hand when it comes to this other kind of charity.

#### 4.2.2 *The Linguistic Sensitivity Principle of Charity*

Here is another principle of interpretive charity: one ought *ceteris paribus* to construct a semantics for ‘knows’ (and its cognates, and perhaps other important epistemic terms) that interprets speakers’ knowledge-attributing behavior in such a way that it regards speakers as sensitive to the contribution ‘knows’ (etc.) makes to the truth-conditions of a proposition using the term when the speaker understands all of the other terms in the sentence. Much less precisely, the principle says we ought to try to interpret others’

verbal behavior so as to regard them as understanding the words they use. It is the presumption of linguistic sensitivity discussed earlier, now in the form of a principle of theory selection.

Let us call this the linguistic sensitivity principle of charity. Just as the truth-maximization principle of charity couldn't tell us whether a proposed semantics was flat-out charitable with respect to truth-maximization, it is unclear whether there is a decisive standard for declaring a proposed semantics flat-out charitable in the linguistic sensitivity sense. As before, comparative judgments, especially concerning the *ceteris paribus* clause, must be made.

The intuitive motivation for this principle of charity should be clear enough. The basic idea is that it is charitable in this sense to regard speakers as 'knowing' what their words mean, including the word 'knows' (etc.). We presuppose that speakers competent with a language are sensitive to the semantics and pragmatics of familiar terms in their language.<sup>118</sup> Indeed, I am inclined to think that sentence is analytically true: competence with a language *just is* sensitivity to the semantics and pragmatics of familiar terms in that language. As we saw, Putnam pointed out long ago that a group of language-users counts one of its members as semantically competent only if that speaker accepts as true certain sentences and rejects as false certain others (1975, p. 250).<sup>119</sup> Someone counts as

---

<sup>118</sup> This should not be taken to imply that ordinary speakers can *explain* the difference between semantic and pragmatic content of utterances, just that they are sensitive to both.

<sup>119</sup> I have received the objection that semantic externalism is incompatible with the presumption of linguistic sensitivity. That is what motivates the use of Putnam here. Putnam staunchly defended the externalist thesis "meaning is reference" and still required semantic competence in the form of what I am calling linguistic sensitivity. It is hard to imagine how any communication could be possible if speakers were not sensitive to the correct behavior of terms familiar to them in the way described here.

“knowing the meaning” of ‘tiger’ only if that person accepts that stereotypical tigers are striped (or, in terminology closer to Putnam’s, accepts that “tigers are striped” is true). Someone would count as not “knowing the meaning” of ‘tiger’ if that person were to say, “That’s a tiger!” and be inclined to retract that claim if someone else responded, “There are ants nearby.” That is because a speaker competent with ‘tiger’ is sensitive to what is and is not relevant to the correct use of ‘tiger,’ and the presence of ants is irrelevant to whether or not there are tigers.<sup>120</sup> The point is that we take for granted that competent speakers accept certain claims and reject certain others, and their tendencies in that regard are a necessary condition of semantic competence. Since ‘knows’ and “S knows that p” are familiar expressions in a language with which, we presume, many speakers are generally semantically competent, the linguistic sensitivity principle of charity would have us interpret their uses of those expressions in such a way as to regard them as semantically competent with those expressions.

The linguistic sensitivity principle of charity easily handles false speech in the form of exaggeration and loose talk. The speaker who uses ‘a million’ to convey “some indeterminately large quantity” is open to challenge by an interlocutor asking, “You don’t really mean ‘a million’, right?” The speaker who understands ‘a million’ is in a position to recognize that the initial claim was false, and can readily respond with, “There were a lot of them,” or “You know what I mean,” or a more precise specification of how many there were. That is because the speaker is sensitive to the fact that the expression ‘a

---

<sup>120</sup> *Semantically* irrelevant, that is. It is a contingent fact that tigers live where there are ants. If this were not so, then ‘There are ants!’ could be taken to challenge ‘That’s a tiger.’ But that possible world is not ours.

million' does *not* mean "some indeterminately large quantity": the speaker was sensitive to that fact about 'a million' *prior* to making the false utterance, and that is why the speaker so quickly recognized that the interlocutor's challenge was relevant to the first claim ("There were a million," or whatever). That the speaker is sensitive to the contribution 'a million' makes to the truth-conditions of a proposition of which it is a part explains this behavior: in other words, the hypothesis that the speaker is sensitive to the behavior of 'a million' predicts this behavior.

Similar remarks apply to loose talk. That the speaker in the coffee example understands what it would take for there to literally be no coffee predicts that, were he to falsely assert that there is no coffee while being simultaneously aware that, in fact, there is some coffee, he would retract or qualify his previous claim under certain conditions.

Neither of these explanations requires a speaker who understands an expression to have an exhaustive list of truth-conditions for the expression fully consciously present to their minds while speaking. But when a speaker is sensitive to the meaning of an expression or term, the speaker will respond in certain predictable ways when the speaker understands all of the other terms in the utterance and is at least disposed to become aware that the utterance made is false because of the use of the term or expression in question.

This idea is borne out when a person sensitive to the meaning of a term encounters an argument which challenges the person's accepted definition of the term. Suppose D is an incorrect definition of X. A person might consciously believe that D is the correct definition of X in the sense that if he were asked for the definition of X, he

would say “D.” The person is still often sensitive to the meaning of X in some way. Sometimes, a convincing argument will lead a person who is sensitive to the meaning of a term to reject as incorrect a definition the person previously accepted. The ability to understand and sometimes acknowledge counterexamples and counterarguments relevant to the definition of a term *presupposes* some kind of grasp of the meaning of that term. For example, there has been a little controversy over how widely philosophers prior to Gettier (1963) accepted the fallibly justified true belief analysis of knowledge. But it is uncontroversial that after Gettier’s paper hardly anyone wanted to defend that account of knowledge. The fact that philosophers were so quick to reject the account suggests that they were primed to become aware that the fallibly justified true belief account is not the correct account of ‘knowledge.’ They were sensitive to the semantics of ‘knowledge’ beforehand.<sup>121</sup>

When a subject understands a term in a language, a set of subjunctive conditionals are true. The antecedents of these conditionals relate the subject, the term, and a scenario wherein the subject considers some state of affairs and whether it would be correctly described as an instance of that term. The conditionals describe the behavior of the subject under those conditions. They have the form, “If S were to consider whether x is F, then S would y.” For example, “If Chisholm were to consider whether Smith truly-but-luckily believing ‘the man with 10 coins in his pocket gets the job’ is an instance of

---

<sup>121</sup> I have tried to state this in a way that is not overly cognitive. One need not presuppose that speakers have dispositional *beliefs* about definitions to accept this account. It is consistent with this account that one’s understanding of a term amounts to no more than dispositions to say certain things in response to various stimuli (prior thoughts, expressions from others, etc.).

knowledge, then Chisholm would deny it.” What makes these subjunctives true is that the subject has a disposition to accept as correct certain descriptions of hypothetical states of affairs, and to reject as incorrect certain descriptions of others.<sup>122</sup> On the proposed account, these dispositions are what understanding the meanings of the terms in one’s own language consists.<sup>123</sup>

If that or something approximately like it is correct, then a proposed semantics for ‘knows’ is charitable in the current sense to the degree that it accurately predicts that speakers will discourse as though they have the relevant dispositions: which, as suggested, will be manifested through verbal behavior in accordance with the subjunctive conditionals made true by the speaker’s dispositions. When a semantic theory says ‘skyscraper’ excludes ‘tree,’ and a particular speaker is familiar with both trees and skyscrapers, then the semantic theory is charitable to the speaker when it anticipates that were S to consider whether a skyscraper is correctly (literally) described by ‘tree,’ S would deny it. The theory in question is charitable because it predicts that S’s verbal behavior will be that of a speaker who *understands* some important boundaries of ‘tree.’

---

<sup>122</sup> It is a stretch to say that the existence of a dispositional property in some entity is *sufficient* to make true subjunctive conditionals describing the behavior of that entity under various circumstances. Something could always interfere so as to prevent the disposition from exercising. In that case, the entity would possess the disposition, but that would fail to make true the subjunctive conditional(s) in question. I want to set that issue aside here. For discussion, see (Chisholm 1955, Lewis 1997). I am committed, though, to the idea that some analysis of these subjunctives in terms of dispositions is part of the correct analysis. Surely *something* makes the subjunctives true, and there are no natural candidates besides dispositions, even if the dispositions themselves are not sufficient.

<sup>123</sup> The account described in this paragraph has much in common with Fumerton (1983). In offering his account of philosophical analysis, Fumerton does not discuss the relevance of scenarios wherein an assertion is challenged, but he does seem to equate understanding an expression with linguistic dispositions of a certain sort.

When a term or expression is being used to assert something false, the speaker who makes the false assertion and understands the term will be open to correction: that speaker will have a disposition making true subjunctive conditionals describing the circumstances under which the speaker would deny that the term or expression is a correct description of the state of affairs in question. If a speaker who uses ‘a million’ to exaggerate understands what ‘a million’ means, then it will be true that if the speaker were challenged on that exaggerated use of ‘a million’, the speaker would understand that the challenge is relevant to the truth of the proposition the utterance expressed, and then would retract the claim, or offer a concession of some sort.

Now for the comparative assessment: how do our candidate views fare with respect to the linguistic sensitivity principle of charity? It should be clear enough that if contextualism is true, the fact that speakers do not generally respond in a way that contextualism predicts they should respond is a mark against contextualism as far as the linguistic sensitivity version of the principle of charity is concerned.<sup>124</sup> If speakers understand ‘knows’ the contextualist way, then they *should* be disposed to respond to the raising of an error possibility by an interlocutor, whether an ordinary possibility or challenge (“Are you *sure* you parked the car in the North part of the lot and not the South?”), or a skeptical one (“Are you *sure* you’re not a car-less brain-in-a-vat?”), by

---

<sup>124</sup> DeRose (2006) offers many arguments for the acceptability of contextualism in the light of Schiffer’s (1996), Hawthorne’s (2004), and Stanley’s (2005) semantic blindness objections. Responding to them all is outside the scope of this paper. But I will address the main theme of his paper: his argument seems to be that because “well-constructed” cases often elicit the intuition that speakers are not contradicting each other when they utter surface-contradictory knowledge attributions suggests that contextualism does not uniquely have a problem of semantic blindness. That may be, but it leaves untouched both that other explanations of the intuition are possible and that contextualism is *far* less intuitively acceptable to most people than an contextual interpretation of “flat”, as Cohen has noted (2004).

pointing out that those challenges are generally not relevant to whether or not the original claim to know was true. Of course, that is not how people respond.

In this paper, I do not want to enter the fray on which version of the semantic blindness objection to contextualism is the most likely to stick. Others have made several cases, at least a few of which have been conceded by contextualists as ways in which speakers are ignorant of language.<sup>125</sup> Here I just indicate that if anything like the contextualist account is correct, speakers ought to reply to challenges to knowledge claims differently than they in fact do. Contextualism is uncharitable on this score because the contextualist explanation of a number of cases requires that speakers are unaware of the propositions their knowledge attributions express. By the current definition, that's linguistic insensitivity, and consequently uncharitable.

What about strict moderate invariantism? Strict moderate invariantists argue that many knowledge denials made after the raising of error possibilities are false: that is, the subject in question often still knows the proposition that subject originally claimed to know, even when an error possibility has been raised which the subject cannot eliminate.<sup>126</sup> Insofar as that is the case, those views are as equally uncharitable as contextualism. Subjects generally retract or somehow qualify knowledge attributions when an error possibility has been raised that the subject cannot eliminate. Contextualists account for that in a counterintuitive way: by rendering the knowledge denial true, but compatible with the truth of the original claim. Moderate invariantists who claim the

---

<sup>125</sup> See especially (Cohen 2004, DeRose 2006, Neta 2003).

<sup>126</sup> Cf. Rysiew (2001), (2007), Williamson (2005), Bach (2005), Black (2005), Brown (2006), (Nagel 2008, 2010), (Gerken 2011, 2012).

subject still knows the original proposition have to account for it by claiming that the subject is confused about what is and is not relevant to whether the original proposition is known. This kind of moderate invariantism predicts that subjects will have a disposition to respond to challenges to knowledge attributions by pointing out that the challenge is irrelevant to whether the proposition is known: *unless*, that is, the view allows that speakers are insensitive to the contribution ‘knows’ makes to the content of the knowledge attribution.<sup>127</sup> It is because the speaker regards ‘knows’ as a correct description of the speaker’s relation to the proposition *only if* the error possibility that has been raised can be ruled out by the subject that the speaker says the subject does not know. So strict moderate invariantism predicts that speakers will have a disposition to respond to challenges in a way *contrary* to how people actually respond. That is uncharitable according to the linguistic sensitivity version of the principle of charity.

How about skeptical invariantism? If skeptical invariantism is correct and people are implicitly aware of the meaning of ‘knows,’ then we should expect just the behavior we in fact see when a knowledge claim is met by a challenge. The skeptic’s proposed explanation is that the concept of knowledge speakers actually possess requires eliminating all error possibilities that might be raised. If that is the standard of knowledge people actually accept, and speakers understand ‘knows’ to require that standard to be met, then the skeptical account predicts speakers will have a disposition to

---

<sup>127</sup> Some strict moderate invariantists are basically prepared to grant this, going so far as to offer purely psychological explanations of why speakers make the mistake of regarding knowledge-irrelevant possibilities as undermining knowledge attributions. See (Gerken 2011, 2012, Nagel 2008, 2010). As I see it, their arguments of Nagel and Gerken amount to damage control, not ways out of the objection that if they are right, people are simply incompetent with ‘knows.’

retract or somehow modify knowledge claims made when error possibilities are raised. And that *is* what we generally see. So skeptical invariantism is more charitable in the linguistic sensitivity way than the other accounts considered here.

It is taken for granted by all parties that any proposed analysis of knowledge will answer to the notion of knowledge that normal people actually have. The goal is not to replace the ordinary concept of knowledge with an alternative concept called by the same name, or to suggest or to command that people speak in some novel way. It should be assumed, then, that a proposed account of ‘knows’ will predict that people understand ‘knows’ in a way harmonious with what that account says ‘knows’ means. It would be a struggle to defend an account of ‘cup’ that says only tables are correctly described as ‘cups’ and speakers are massively ignorant of that fact and therefore very resistant to accept the table theory of ‘cup.’ What is defective about that account of ‘cup’ is that it clearly does not answer to how ‘cup’ is understood by competent speakers.

What we have seen is that, given the way people actually talk, our two principles of charity impose contrary obligations upon us. On the truth-maximization principle of charity, the non-skeptical views predict that people use ‘knows’ in a way friendly to their accounts and skepticism does not. But on the linguistic sensitivity principle, the non-skeptical views make predictions unfriendly to their own accounts and skepticism comes out ahead.

It is at least a little tempting at this point to call a draw and say that whichever way of being charitable you prefer will have its pros and cons, so the choice is arbitrary: if you prefer having people be ‘knowledgeable’ about their own language, pick

skepticism; if you prefer having people make true claims more often than not, pick one of the other views.

But there is some reason to think respecting the linguistic sensitivity principle of charity should be preferred to the truth-maximization one. We have already seen that expressions which are widely used to make false claims when the speaker is at least implicitly aware that the claim is false cause difficulty for applying the truth-maximization principle of charity. We encountered no similar problem for the linguistic sensitivity version of the principle.

On the other hand, one might think there is surely something defective about a principle of *charity* that is compatible with people speaking falsely so frequently. There are two things to say here. First, if that objection amounts to claiming that the linguistic sensitivity principle of charity is somehow defective because it is not the truth-maximization principle of charity, that is beside the point (and question-begging in the bad way). We have already seen that the linguistic sensitivity principle of charity deserves the name ‘charity.’ Second, the prospect of a skeptical invariantist semantics respecting the linguistic sensitivity principle of charity rests upon an account of linguistic understanding with the result I argued for above: that speakers are in a position to become aware of the falsity of their knowledge attributions. If that is not the case—if speakers are “semantically blind” to the falsity of most knowledge attributions—then skeptical invariantism *does not* respect the linguistic sensitivity principle of charity better than any of the other positions in the debate.

While some philosophers seem to regard semantic blindness as simply inevitable and therefore more-or-less acceptable (DeRose 2006), I think semantic blindness undermines the very possibility of armchair philosophy. If semantic blindness is widespread and uncorrectable, most epistemologists are in very bad shape when it comes to analyzing ‘knows.’ Most of us have no lexicographical training in the empirical study of word use. *We* are the ordinary speakers who are confused about ‘knows’: confused about the truth-value of common knowledge attributions, confused about whether or not standards for ‘knows’ vary by context (as ‘flat’ and ‘tall’ supposedly do), and so on. How, then, can we accept that speakers are regularly confused about knowledge-talk, and then go on to provide an account of it ourselves?<sup>128</sup> A traditional conception of philosophical analysis understands the project to consist basically in entertaining hypothetical states of affairs and introspecting on whether one is inclined to predicate some term of philosophical interest to that state of affairs. On this conception, if I want to analyze ‘knowledge,’ I run through some thought experiments, see where I am and am not inclined to say the cases are instances of knowledge, and attempt to figure out what the cases of knowledge have in common with each other that is not shared with the cases of non-knowledge. But all of that presupposes that I am competent with ‘knows.’ If I am not—if I am semantically blind with respect to ‘knows’—then I cannot use the traditional thought experiment method. For that reason, I conclude that considerations of interpretive charity favor a skeptical treatment of most ordinary uses of knowledge-talk.

---

<sup>128</sup> It is possible that defenders of the non-skeptical views discussed in this section do not take themselves to be analyzing “S knows that p” the way ordinary speakers would understand that expression. But if that is so, then I do not know what they are doing.

Again, this is not a knock-down argument. But the fact that a skeptical analysis of knowledge predicts patterns of usage consistent with the analysis while the other views do not is a substantial mark in its favor.

#### 4.2.3 *Implicatures of “S Knows That P”*

Skeptical invariantists will say that when someone says “S knows that p,” what is literally expressed is false. According to the analysis of knowledge I have defended in this dissertation, knowledge attributions express that S is directly acquainted with p and with everything that makes p true. Because the proposition expressed by a knowledge attribution is nearly always false, a little more needs to be said about what is conveyed by a knowledge attribution: for it would be strange to go on making knowledge attributions if their use does not even accomplish anything. I have already said that knowledge attributions are used loosely and that speakers are in a position to easily come to accept that this is so: but I have not said much about what knowledge attributions are used to convey. I will do that now.

There are a lot of things one can convey by “S knows that p.” Indeed, a skeptical invariantist can plausibly incorporate a lot of the things other epistemologists identify with the meaning of “S knows that p” into his account using the implicatures of “S knows that p”. Such implicatures will often be true even while what is literally expressed by a knowledge attribution is false. Some epistemologists have argued that a central role of the concept of knowledge is to identify reliable informants (Craig 1990). Others, that one knows that p if and only if one’s basis for believing that p rules out all relevant not-p alternatives (Dretske 1970, 1972). These and similar ideas about what knowledge is can

be put to work by the skeptical invariantist for an account of why we use knowledge-talk even while we are in a position to recognize its widespread falsity.

Suppose someone is looking for a tutor for an upcoming physics exam. A non-expert might plausibly say, “You’re looking for a physics tutor? Ask Phyllis—she knows that subject really well.”<sup>129</sup> What is literally stated—that of the many truths of physics, Phyllis is directly acquainted with those propositions and what makes them true—is false. Still, in a context like this, attributing lots of knowledge of the truths of physics to Phyllis plays a relevant practical role in the context, namely to point the student in need of a tutor in the direction of someone who is a reliable informant about the relevant truths. On the plausible assumption that in this exchange what the speaker is trying to get across to the student seeking a tutor is *that Phyllis is a reliable informant on the matter of physics*, then the implicature will be true. The example is intended to generalize: often, when someone says, “S knows that p” in a particular context, it is false that S knows that p, but it will often be true that S is a reliable informant on the subject-matter salient to the conversation.<sup>130</sup>

Similarly, there is often a rough understanding of what will and will not count as relevant (*for practical purposes*) to the truth-value of a knowledge attribution in a particular context. I do not think a reliable alternatives theory of knowledge is adequate

---

<sup>129</sup> For the sake of making the example sound natural I phrased it so what Phyllis is said to know is physics, a field of study: in order to be true, presumably that would have to mean something like, “There are many truths of physics that may be relevant to the upcoming exam, and Phyllis knows many of them.”

<sup>130</sup> If being a reliable informant is ultimately analyzed in terms of knowledge, then of course the implicature will nearly always be false. But I do not see why being reliable about x should imply knowing many truths about x; my car is a reliable means of transportation, but it does not know anything, nor do I have to know that it will get me from place to place in order to have reasonable confidence in its reliability.

as a theory of knowledge (for reasons outlined in chapter two) or that precise criteria for relevance can be given to spell out a robust relevant alternatives theory of knowledge, but I do think a rough notion of reliable alternatives is often usefully deployed in ordinary talk. Suppose someone is walking through a zoo looking for the zebra exhibit, and stops to ask a passerby: “Do you know where the zebras are?” “Yes, they’re down *that way*, to the right” will be taken to convey that the informant has successfully pointed the way to the zebras *rather than* to tigers or alligators, but not that the informant has successfully pointed the way to the zebras *rather than* cleverly painted mules. The implicature is that the zebras and not some other relevant non-zebra thing are ‘that way and to the right’. Plausibly, that implicature and others similar to it will be true, even on a skeptical invariantist view. Again: what contextualists and moderate invariantists identify as the proposition implicitly expressed by a knowledge attribution, the skeptical invariantist can plausibly identify as an implicature.

The point here is general. There is a lot we do with our knowledge-talk and the utility of such talk can be adequately accounted for by a view that says knowledge-talk is nearly always false provided that there are true propositions conveyed by implicature in conversational contexts that accomplish what speakers intend for their knowledge attributions to accomplish. Such implicatures will often include that someone is a reliable informant on a particular subject-matter, that some state of affairs obtains that can be distinguished from a relevant alternative state of affairs, that some particular proposition is true, that some method of forming beliefs reliably delivers true beliefs, and

one can be counted on to do some particular thing or be some particular place at some particular time, etc.

In this way, knowledge ascriptions are a lot like attributions of large quantities and geometric properties. I have used the example ‘a million’ in several places in this chapter to get across the idea that we often use certain terms loosely. “S knows that p” is nearly always false even though it usually conveys something true; ‘a million’ is often used the same way, and can be used to convey that some quantity was large enough to cause annoyance or feeling overwhelmed or a variety of other things in particular contexts, even though there are clearly not one million of the item in question. The same goes for geometric properties, attributions of which are often used to convey something true even while nearly always expressing a false proposition. If I walk to the whiteboard in front of a class and draw a figure with four ‘sides’ and put ‘boxes’ in each corner, and I ask, “What figure is this?” nearly everyone will say, “It’s a rectangle.” I can get them to say it’s a square if I put tick marks on each ‘line segment’.<sup>131</sup> If I point out that the figure is not even close to being a square, there is a sort of collective shrug: they already knew the figure was woefully far from being a square, but normally such drawings are intended to suggest an abstraction of the figure whose properties will be discussed. Sometimes the ‘lines’ do not connect; no part is straight by even loose standards, etc. But in a context where what is relevant is talking about the properties of squares, inadequate representations like the sort of figure I would draw are tolerated as good enough. If I

---

<sup>131</sup> All the geometric terms received scare quotes here because none of my drawings much resemble shapes with the properties denoted by the terms used to pick out the relevant features of my drawings.

point to the figure I have drawn on the board and I say, “Here is a square,” what I have said is false, and everyone is in a position to acknowledge this. But, in contexts where shape is relevant, my false shape-attribution often still conveys something true. That context may be in geometry class, where a false use of “this is a square” may implicate that the figure suggests enough important properties of ideal squares that it can be used to deduce the length of segment A from the lengths of segments B and C that cut through the figure; or the context may be in a hardware store where a customer is shopping for windows and wants ‘square’ windows in their new home addition, and “these ones are square”, while false, conveys that the windows are visually indistinguishable from actual squares and thus look square. The point, again, is general. We use knowledge attributions in the same way that we use attributions of quantity and shape: to convey that something approximates a certain standard for relevant purposes, even when that standard is not actually met by the thing in question. However, there is utility in making the false claim because of what hearers will take from a knowledge attribution in a particular context.

Some have compared knowledge attributions with exaggeration (Butchvarov 1970, p. 55, Hawthorne 2004, p. 115, Schaffer 2004b).<sup>132</sup> The theory is that when we say “S knows that p,” we do not believe what we say. We say something stronger than what

---

<sup>132</sup> What follows in my description of the exaggeration view is basically shared by the three authors cited. As usual, there are subtle differences. Hawthorne rejects the exaggeration view. Schaffer endorses it. Butchvarov’s presentation is the most detailed, as he gives a reason why speakers who do not accept a standard would go on to make false claims to the effect that the standard is met. He says it is because “the use of such a word is likely to produce in the listener important attitudes and perhaps encourage him to act in certain ways” (p. 55). On his view, knowledge attributions (and other claims that are widely false as a result of their demanding standards of application) have a practical purpose that is accomplished by being made even though false.

we believe to be the case. Arguably, much of the ordinary use of “a million” that I have pointed to is exaggeration. We are not even inclined to believe there were a million of whatever. On the exaggeration view, “I know that p” or “So-and-so knows that p” falsely states but truly implicates that some important range of possibilities can be eliminated, even though no one making the claims sincerely believes *all* of the possibilities can be eliminated.

There is an important difference between the loose talk account and the exaggeration account. The exaggeration account adds to the loose talk account the claims that (1) speakers do not believe their knowledge attributions are true, and (2) knowledge claims are deliberate overstatements of the epistemic position the speaker takes the subject to be in (presumably made for some practical purpose). So, the exaggeration theory is a more specific version of the loose talk theory. Sometimes we do use knowledge attributions as a kind of deliberate exaggeration that we believe is false: “I know you’ll do great!” before a friend’s performance or talk. We are of course aware that we know no such thing. But knowledge attributions are not merely exaggerations. As general claims about knowledge attributions, (1) and (2) are both false: subjects plainly do often think “S knows that p” is true, and are not nearly so deliberate about making knowledge claims as to overstate a subject’s epistemic position to achieve some communicative purpose. Sometimes speakers are simply trying to assert “S knows that p” to state that S knows that p.

The exaggeration account is too demanding, even though knowledge attributions are sometimes exaggerations. That is why I prefer the loose talk account that says

knowledge attributions are widely false and speakers are in a position to become aware of this. The loose talk account by itself does not say anything about what speakers are trying to accomplish through making knowledge claims. It is important to avoid committing to some particular thesis about what all speakers are always trying to accomplish with knowledge attributions: there are just too many ends for which a knowledge claim can serve as an effective means. Often, if we want to find out what a speaker intends to express through a knowledge claim, we must simply ask: “What do you mean?”

Call an emphatic knowledge attribution a claim of the form “S knows that *p* *for sure*”, “S knows that *p* *for certain*”, “S *really* knows that *p*”, and so on. On the surface, an emphatic knowledge attribution states that S meets some kind of epistemic standard stricter than knowledge, or that S knows with the addition of some property that guarantees S’s knowledge. Some epistemologists identify epistemic concepts whose standard of application is stronger than knowledge, for instance certainty (“epistemic” certainty). As I argued in previous chapters, knowing that *p* requires being acquainted with everything that makes *p* true; as a result, there is no better position one could be in with respect to *p*’s truth. For that reason, there are no epistemic concepts whose standard of application is more demanding than knowledge. Why, then, do people make emphatic knowledge attributions?

An emphatic knowledge attribution might point to the possession of a property by S that one could think is relevant to knowledge, even though the property is not conceptually related to truth. For instance, I might say “I know for sure” intending to

state not only that I know that p but also that I am not at all in doubt regarding the truth of p. So, an emphatic knowledge attribution might state both that S knows and that S has the additional property of being maximally confident. Being confident has no bearing on my acquaintance with the factors that constitute p's being true. So while an emphatic knowledge attribution does not attribute a higher epistemic status to S's belief than a non-emphatic knowledge attribution does, one might use an emphatic knowledge attribution in slightly different contexts than a non-emphatic knowledge attribution. For example, an emphatic knowledge attribution can be effectively used to reassure someone who is in doubt, or to put down a proposed challenge to a non-emphatic knowledge attribution. As with other knowledge attributions, there are many purposes to which one might put an emphatic knowledge attribution. Often, the only way to figure out what purpose a speaker has in mind is to ask.

#### 4.2.4 *Knowledge Attributions in High but Non-Philosophical Contexts*

The skeptic can go quite a long way accounting for the common pattern of using 'knows' to make false claims by means of the usage patterns I identified in this section: exaggeration, loose talk, and conditionalization. But there is undoubtedly one more factor that plays a role in our talk about knowledge: epistemological error. If skepticism is true, then when someone says, "S knows that p," and the skeptic replies by raising various error-possibilities not known by the speaker to be eliminated or eliminable by S, if that speaker continues to assert "S knows that p," the speaker must simply be mistaken. The speaker is misevaluating the strength of epistemic position occupied by the subject of the knowledge attribution—a blatant epistemological error.

Epistemological error is unlike the kinds of error the other views in this debate have to worry about because epistemological error is not an error with respect to a speaker's semantic competence, but an error concerning the facts of a particular case or class of cases. Still, it does not follow that epistemological error is not just as bad or worse than semantic incompetence. Consequently, for the sake of the skeptic's case, the less epistemological error, the better.

Where might we find cases of uncorrectable epistemological error? I suspect cases in which speakers are in raised epistemic contexts where the standards are raised not by general skeptical scenarios, like en-vatting or radical deception by a powerful malevolent being, but by consideration of possible mistakes in practical contexts under raised epistemic standards. To the surgeon: "Do you *know* whether or not you got all of the medical tools out of the patient's belly, before we stitch him up?" To the trial witness: "Can you point to the person who shot your friend that night? Are you *sure* that's him?" To the distressed teacher: "Do you know if all of the kids were in from the playground before you locked the door during the tornado?" In these scenarios, there are practical stakes high enough that knowledge attributions are not made casually and are often challenged. Still, the skeptic has to say these knowledge attributions are false, too: and if the speakers are not disposed to retract, conditionalize, or somehow clarify, the skeptic has to chalk them up to cases of epistemological error.

Are high but not skeptically-high contexts ones where, contrary to the normal pattern, the skeptic's moves to get speakers to recognize the falsity of knowledge claims fail? That would be a somewhat odd result. While there is some intuitive pull to say

speakers under the kind of serious pressure in the above scenarios are not speaking loosely, the skeptic will insist that, if given a few minutes with the speakers where they did not have to worry about the legal or further moral implications of anything they are about to say being heard, they would admit under skeptical questioning that ‘knows’ is just not the right way to describe the epistemic relation they bear to the claims under consideration. There might be serious practical constraints on getting the speakers to acknowledge this, but all along we have not been talking about what is *easy* to get people to say, but what they most likely would say in response to a line of questioning singly focused on the perceived truth-values of epistemic claims.<sup>133</sup>

Finally, it must be noted that I have not explicitly identified what “skeptical standards” are in this section, except in saying that skeptical invariantism will claim that knowledge attributions are usually or always false. I stand by the epistemological views advanced in the previous chapters. So I think when someone says “S knows that p” the proposition expressed says S is directly acquainted with what makes p true. But the argument in defense of skeptical invariantism about knowledge attributions is designed to cover not only my preferred analysis of knowledge, but other skeptical accounts. Well, the more severe the skepticism, the more error or semantic incompetence that has to be attributed by the skeptic to speakers. If some version of skepticism says no one has any epistemic justification for favoring any proposition over its contrary, that view will face

---

<sup>133</sup> For partly this reason, the fact that experimental philosophers poll undergraduates who say this-or-that in response to some prompt should be taken note of, but is merely a single piece of evidence for any philosophical theory: it is just too hard with such a population to control for the kinds of things trained philosophers pay attention to without needing to be instructed.

an increasingly steep uphill battle in accommodating ordinary talk. It is one thing to defend high standards for knowledge against the other views we have considered here: in that case, one might argue that while no one knows anything, we very often have good justification for many of the things we believe. That kind of skepticism is, I believe, comparatively not too difficult to motivate. And while the justification skeptic gets some traction by the patterns of conditionalization, withdrawal, exaggeration, and acknowledgement of loose use for knowledge attributions, those patterns will not easily go as far in accommodating what people think they justifiably believe: which is a lot, even upon some questioning. It is crucial to see here that the same patterns of use in response to various questions and objections in what we think of as ‘ordinary’ discourse is *structurally* exactly the same as what goes on in epistemology class. The fact that it takes a little *more* pushing to get someone to accept that the claim that justification might be very hard to come by is no objection either to the truth of that severe skeptical hypothesis or to the notion that speakers understand “justified” and “rational” in such a way that they are sensitive to skeptical challenges.

### 4.3 Chapter Summary

In this section, I have argued that the methodology argued for early in the dissertation supports accepting the linguistic sensitivity principle of charity: a principle that favors attributing semantic competence concerning ‘knows’ over maximizing the truth-values of classes of knowledge attributions. Considerations of semantic competence, in turn, favor skepticism. It is hard to make a comparative evaluation of why it is better to favor semantic competence over truth-maximization except that favoring semantic competence

for epistemic terms fits better with our argued-for methodology and neglecting semantic competence threatens the possibility of armchair philosophy. I pointed to a few patterns of speech and indicated that the skeptic can assimilate knowledge claims to those uses of language without requiring much in the way of semantic incompetence or epistemological error. There is still a conditional worry, though: to whatever degree speakers cannot be brought to give answers to lines of questioning in a way consistent with what the skeptic identifies as the correct epistemic standards, to that degree speakers are semantically incompetent or victims of epistemological error. And the more those problems abound, the more skepticism fails to best the other views on the issue of knowledge attributions.

## AFTERWORD

In the epigraph, Descartes expressed that he could not go too far with his attitude of methodical doubt because his inquiry concerned only knowledge and not action. I have for a long time had little difficulty accepting that human knowledge is very limited. However, perhaps the dominant opinion in epistemology today is that knowledge and action cannot be so easily pulled apart. Knowledge has been identified as requirement of everything from asserting to acting on a reason. For instance, if someone asserts that *p*, an interlocutor may ask, “How do you know that *p*?” The question seems legitimate. If the speaker cannot provide an account of how *p* is known, we take it that the speaker has done something worthy of criticism. One should not just go around making assertions if one does not *know* that what one says is true. Similarly, suppose someone reasons as follows (prior to the results being announced): ‘This lottery ticket is a loser. So, if I keep it, I will gain nothing. If I sell it, I will get a penny. So, I should sell it’ (Hawthorne 2004). That reasoning seems suspicious. Many philosophers think the suspicion is warranted because one is not entitled to the premise ‘this ticket is a loser’ because the truth of that proposition is not known.

The point is that there may be important connections between knowledge and warranted assertion and warranted action. Whatever anyone says about the plausibility of skepticism when it is limited to the purely epistemic domain, there is a real concern that skepticism may have severe implications for our talk and action. That is to say there may be no *purely* epistemic domain if the correct analysis of non-epistemic concepts ultimately involves epistemic concepts. It seems to me that there are really only two

ways out. One is to deny that knowledge has these conceptual connections, so the apparent legitimacy of criticism for not knowing a proposition one asserts or for relying on a premise in practical reasoning when one does not know that premise merely appear legitimate but are not. The other option is to accept that knowledge has these conceptual connections with assertion, practical reasoning, and perhaps more, but come up with a plausible theory of excuses: though there may be something wrong with asserting that p and not knowing that p, it may just be that the knowledge norm of assertion is better violated than not. The same may go for practical reasoning. I am not sure which way to go, but my plans for a comprehensive, plausible skeptical analysis of knowledge include addressing the issue.

In any case, I've argued that a wide-ranging skepticism is plausible. The remaining issue is whether or not we can live with it.

## REFERENCES

- Abath, A. J. (2012). "Epistemic Contextualism, Semantic Blindness and Content Unawareness." Australasian Journal of Philosophy **90**(3): 593-597.
- Alston, W. (1989). Level Confusions in Epistemology. Epistemic Justification: Essays in the Theory of Knowledge. Ithaca, Cornell University Press: 153-171.
- Armstrong, D. M. (1973). Belief, Truth, and Knowledge. New York, Cambridge University Press.
- Audi, R. (1993). The Structure of Justification, Cambridge University Press.
- Ayer, A. J. (1956). The Problem of Knowledge. London, Macmillan.
- Bach, K. (2005). The Emperor's New 'Knows'. Contextualism in philosophy: Knowledge, meaning, and truth. G. Preyer and G. Peter. New York, Oxford University Press: 51-89.
- Bergmann, M. (2006). Justification without Awareness. New York, Oxford University Press.
- Black, T. (2005). "Classic Invariantism, Relevance and Warranted Assertability Maneuvres." The Philosophical Quarterly **55**(219): 328-336.
- BonJour, L. (1985). The Structure of Empirical Knowledge. Cambridge, Harvard University Press.
- BonJour, L. (1998). In Defense of Pure Reason. New York, Cambridge University Press.
- BonJour, L. (2010). "The Myth of Knowledge." Philosophical Perspectives **24**(1): 57-83.
- BonJour, L. and E. Sosa (2003). Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues. New York, Oxford University Press.

- Bourget, D. and D. J. Chalmers (2014). "What Do Philosophers Believe?" Philosophical Studies **170**: 465-500.
- Brown, J. (2005). "Comparing Contextualism and Invariantism on the Correctness of Contextualist Intuitions." Grazer Philosophische Studien **69**(1): 71-100.
- Brown, J. (2006). "Contextualism and Warranted Assertibility Manoeuvres." Philosophical Studies **130**(3): 407-435.
- Burge, T. (1979). "Individualism and the Mental." Midwest Studies in Philosophy **4**(1): 73-121.
- Butchvarov, P. (1970). The Concept of Knowledge. Evanston, Northwestern University Press.
- Carnap, R. (1955). "Meaning and Synonymy in Natural Languages." Philosophical Studies **6**(3): 33-47.
- Chalmers, D. J. (1996). The Conscious Mind. New York, Oxford University Press.
- Chalmers, D. J. (2003). "The Nature of Narrow Content." Philosophical Issues **13**(1): 46-66.
- Chisholm, R. M. (1955). "Law Statements and Counterfactual Inference." Analysis **15**(5): 97 - 105.
- Chisholm, R. M. (1977). Theory of Knowledge, Second Edition. Englewood Cliffs, Prentice-Hall, Inc.
- Chrisman, M. (2007). "From Epistemic Contextualism to Epistemic Expressivism." Philosophical Studies(135): 225-254.

- Clark, M. (1963). "Knowledge and Grounds: A Comment on Mr. Gettier's Paper." Analysis **24**(2): 46-48.
- Cohen, S. (1984). "Justification and Truth." Philosophical Studies **46**(3): 279-295.
- Cohen, S. (1986). "Knowledge and Context." The Journal of Philosophy **83**(10): 574-583.
- Cohen, S. (1988). "How to be a Fallibilist." Philosophical Perspectives **2**: 91-123.
- Cohen, S. (1998). "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." Australasian Journal of Philosophy **76**(2): 289-306.
- Cohen, S. (1999). "Contextualism, Skepticism, and the Structure of Reasons." Noûs **33**: 57-89.
- Cohen, S. (2004). "Contextualism and Unhappy-Face Solutions: Reply to Schiffer." Philosophical Studies **119**(1/2): 185-197.
- Conee, E. (2007). Externally Enhanced Internalism. Internalism and Externalism in Semantics and Epistemology. S. C. Goldberg. New York, Oxford University Press: 51-67.
- Conee, E. and R. Feldman (2001). Internalism Defended. Epistemology: Internalism and Externalism. H. Kornblith. Oxford, Blackwell: 231-260.
- Conee, E. and R. Feldman (2004). Evidentialism: Essays in Epistemology. New York, Oxford University Press.
- Craig, E. (1990). Knowledge and the State of Nature. Oxford, Clarendon Press.
- Davidson, D. (1973). "Radical Interpretation." Dialectica **27**(3-4): 313-328.
- Davis, W. A. (2006). "Knowledge Claims and Context: Loose Use." Philosophical Studies **132**(3): 395-438.

- Davis, W. A. (2013). "On Nonindexical Contextualism." Philosophical Studies **163**(2): 561-574.
- DeRose, K. (1992). "Contextualism and Knowledge Attributions." Philosophy and Phenomenological Research **52**(4): 913-929.
- DeRose, K. (1995). "Solving the Skeptical Problem." The Philosophical Review **104**(1): 1-52.
- DeRose, K. (2005). "The Ordinary Language Basis for Contextualism, and the New Invariantism." The Philosophical Quarterly **55**(219): 172-198.
- DeRose, K. (2006). "'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments against Contextualism." Philosophy and Phenomenological Research **73**(2): 316-338.
- DeRose, K. (2008). "Gradable Adjectives: A Defence of Pluralism." Australasian Journal of Philosophy **86**(1): 141-160.
- Dictionary, O. E. "knowledge, n.", Oxford University Press.
- Dougherty, T. (2011). In Defense of Propositionalism about Evidence. Evidentialism and its Discontents. T. Dougherty. New York, Oxford University Press: 226-232.
- Dougherty, T. and P. Rysiew (2009). "Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions." Philosophy and Phenomenological Research **78**(1): 123-132.
- Dougherty, T. and P. Rysiew (2011). "Clarity about Concessive Knowledge Attributions: Reply to Dodd." Synthese **181**(3): 395-403.

- Dretske, F. (1971). "Conclusive Reasons." Australasian Journal of Philosophy **49**(1): 1-22.
- Dretske, F. (1981). "The Pragmatic Dimension of Knowledge." Philosophical Studies **40**(3): 363-378.
- Dretske, F. I. (1970). "Epistemic Operators." The Journal of Philosophy **67**(24): 1007-1023.
- Dretske, F. I. (1972). "Contrastive Statements." The Philosophical Review **81**(4): 411-437.
- Fales, E. (1996). A Defense of the Given. Lanham, Rowman and Littlefield.
- Fantl, J. and M. McGrath (2009). Knowledge in an Uncertain World. New York, Oxford University Press.
- Feldman, R. (2004). Having Evidence. Evidentialism: Essays in Epistemology. E. Conee and R. Feldman. New York, Oxford University Press: 219-241.
- Feldman, R. and E. Conee (1985). "Evidentialism." Philosophical Studies **48**: 15-34.  
Reprinted in Evidentialism: Essays in Epistemology, Oxford University Press (2004).
- Fine, K. (2002). The Varieties of Necessity. Conceivability and Possibility. T. S. Gendler and J. Hawthorne. New York, Oxford University Press.
- Fodor, J. and E. LePore (1992). Introduction: A Geography of the Issues. Holism: A Shopper's Guide. J. Fodor and E. LePore. Cambridge, Blackwell Publishers.
- Frege, G. (1884 (1997)). The Foundations of Arithmetic. Malden, Blackwell Publishing.

- Fumerton, R. (2006a). Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament. Knowledge and Reality: Essays in Honor of Alvin Plantinga. T. Crisp, M. Davidson and D. Vander Laan, Springer Netherlands. **103**: 179-191.
- Fumerton, R. (2006b). Epistemology. Malden, Wiley-Blackwell.
- Fumerton, R. (2010). "Poston on Similarity and Acquaintance." Philosophical Studies **147**(3): 379-386.
- Fumerton, R. A. (1983). "The Paradox of Analysis." Philosophy and Phenomenological Research **43**(4): 477-497.
- Fumerton, R. A. (1985). Metaphysical and Epistemological Problems of Perception. Lincoln, University of Nebraska Press.
- Fumerton, R. A. (1995). Metaepistemology and Skepticism, Rowman & Littlefield Publishers.
- Gendler, T. S. and J. Hawthorne, Eds. (2002). Conceivability and Possibility. New York, Oxford University Press.
- Gerken, M. (2011). "Epistemic Focal Bias." Australasian Journal of Philosophy **91**(1): 41-61.
- Gerken, M. (2012). On the Cognitive Bases of Knowledge Ascriptions. Knowledge Ascriptions. J. Brown and M. Gerken. New York, Oxford University Press: 140-170.
- Gettier, E. L. (1963). "Is Justified True Belief Knowledge?" Analysis **23**(6): 121-123.
- Goldman, A. and J. Pust (1998). Philosophical Theory and Intuitional Evidence. Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical

- Inquiry. M. R. DePaul and W. Ramsey. Lanham, Rowman & Littlefield Publishers, Inc.: 179-197.
- Goldman, A. I. (1967). "A Causal Theory of Knowing." The Journal of Philosophy **64**(12): 357-372.
- Goldman, A. I. (1976). "Discrimination and Perceptual Knowledge." The Journal of Philosophy **73**(20): 771-791.
- Goodman, N. (1955). Fact, Fiction, and Forecast. Cambridge, Harvard University Press.
- Grice, H. P. (1957). "Meaning." Philosophical Review(66): 377-388.
- Grice, H. P. (1975). Logic and Conversation. Syntax and Semantics, vol. 3: Speech Acts. P. Cole and J. Morgan. New York, Academic Press.
- Hare, R. M. (1960). "Philosophical Discoveries." Mind: A Quarterly Review of Philosophy **LXIX**(274): 145-162.
- Harman, G. (1973). Thought. Princeton, Princeton University Press.
- Hawthorne, J. (2004). Knowledge and Lotteries. New York, Oxford University Press.
- Hume, D. (1739 (2000)). A Treatise of Human Nature. New York, Oxford University Press.
- Hume, D. (1777). An Enquiry Concerning Human Understanding.
- Jackson, F. (1998). From Metaphysics to Ethics. New York, Oxford University Press.
- Kant, I. (1781 (2006)). The Critique of Pure Reason. New York, Cambridge University Press.
- Keynes, J. M. (1921). A Treatise on Probability. London, Macmillan & Co.

- Klein, P. (1998). "Foundationalism and the Infinite Regress of Reasons." Philosophy and Phenomenological Research **58**(4): 919-925.
- Korcz, K. A. (2015). "The Epistemic Basing Relation." The Stanford Encyclopedia of Philosophy **Fall 2015 Edition**.
- Kornblith, H. (2002). Knowledge and its Place in Nature. New York, Oxford University Press.
- Kripke, S. (1980). Naming and Necessity. Cambridge, Harvard University Press.
- Lehrer, K. (1971). "How Reasons Give us Knowledge, or the Case of the Gypsy Lawyer." The Journal of Philosophy **68**(10): 311-313.
- Lehrer, K. (1990). Theory of Knowledge. Abingdon, Routledge.
- Lehrer, K. and T. Paxson (1969). "Knowledge: Undefeated Justified True Belief." The Journal of Philosophy **66**(8): 225-237.
- Lewis, C. I. (1946). An Analysis of Knowledge and Valuation. La Salle, Open Court.
- Lewis, D. (1996). "Elusive knowledge." Australasian Journal of Philosophy **74**(4): 549-567.
- Lewis, D. (1997). "Finkish Dispositions." The Philosophical Quarterly **47**(187): 143-158.
- Ludlow, P. (2005). Contextualism and the New Linguistic Turn in Epistemology. Contextualism in philosophy: Knowledge, meaning, and truth. G. Preyer and G. Peter. New York, Oxford University Press: 11-50.
- MacFarlane, J. (2007). "Relativism and Disagreement." Philosophical Studies **132**(1): 17-31.
- MacFarlane, J. (2009). "Nonindexical Contextualism." Synthese **166**(2): 231-250.

- McGrew, T. (1995). The Foundations of Knowledge. Lanham, Rowman and Littlefield Publishers.
- McGrew, T. and L. McGrew (2007). Internalism and Epistemology: The Architecture of Reason. New York, Routledge.
- McKinsey, M. (1991). "Anti-Individualism and Privileged Access." Analysis **51**(1): 9-16.
- McLaughlin, B. P. (2003). McKinsey's Challenge, Warrant Transmission, and Skepticism. New Essays on Semantic Externalism and Self-Knowledge. S. Nuccetelli. Cambridge, MIT Press: 79-96.
- Montminy, M. (2009). "Contextualism, Invariantism and Semantic Blindness." Australasian Journal of Philosophy **87**(4): 639-657.
- Moore, G. E. (1903 (2000)). Principia Ethica. Cambridge, Cambridge University Press.
- Moser, P. K. (1991). Knowledge and Evidence, Cambridge University Press.
- Myers-Schulz, B. and E. Schwitzgebel (2013). "Knowing That P without Believing That P." Noûs **47**(2): 371-384.
- Nagel, J. (2008). "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes." Australasian Journal of Philosophy **86**(2): 279-294.
- Nagel, J. (2010). "Knowledge Ascriptions and the Psychological Consequences of Thinking about Error." The Philosophical Quarterly **60**(239): 286-306.
- Neta, R. (2003). "Skepticism, Contextualism, and Semantic Self-Knowledge." Philosophy and Phenomenological Research **67**(2): 396-411.
- Nozick, R. (1981). Philosophical Explanations, Harvard University Press.

- Pritchard, D. (2002). "Two Forms of Epistemological Contextualism." Grazer Philosophische Studien(64): 19-55.
- Pritchard, D. (2005). Epistemic Luck. New York, Oxford University Press.
- Putnam, H. (1962). "It Ain't Necessarily So." The Journal of Philosophy **59**(22): 658-671.
- Putnam, H. (1973). "Meaning and Reference." The Journal of Philosophy **70**(19): 699-711.
- Putnam, H. (1975). The Meaning of 'Meaning'. Mind, Language, and Reality: Philosophical Papers. New York, Cambridge University Press. **2**.
- Radford, C. (1966). "Knowledge--by Examples." Analysis(27): 1-11.
- Rey, G. (2013) "The Analytic/Synthetic Distinction." The Stanford Encyclopedia of Philosophy (Fall 2013 Edition) **2015**.
- Richard, M. (2004). "Contextualism and Relativism." Philosophical Studies **119**(1-2): 215-242.
- Roush, S. (2005). Tracking Truth: Knowledge, Evidence, and Science. New York, Oxford University Press.
- Roush, S. (2012). Sensitivity and Closure. The Sensitivity Principle in Epistemology. K. Becker and T. Black. New York, Cambridge University Press: 242-268.
- Russell, B. (1905). "On Denoting." Mind: A Quarterly Review of Philosophy **14**(456): 479-493.
- Russell, B. (1910). "Knowledge by Acquaintance and Knowledge by Description." Proceedings of the Aristotelian Society **11**: 108-128.
- Russell, B. (1912). The Problems of Philosophy, Oxford University Press

- Russell, G. (2008). Truth in Virtue of Meaning. New York, Oxford University Press.
- Ryle, G. (1949). The Concept of Mind. Chicago, University of Chicago Press.
- Ryle, G. (1970). Autobiographical. Ryle: A Collection of Critical Essays. O. P. Wood and G. Pitcher. New York, Doubleday.
- Rysiew, P. (2001). "The Context-Sensitivity of Knowledge Attributions." Noûs **35**(4): 477-514.
- Rysiew, P. (2007). "Speaking of Knowing." Noûs **41**(4): 627-662.
- Schaffer, J. (2004a). "From Contextualism to Contrastivism." Philosophical Studies **119**(1/2): 73-103.
- Schaffer, J. (2004b). "Skepticism, Contextualism, and Discrimination." Philosophy and Phenomenological Research **69**(1): 138-155.
- Schaffer, J. (2010). "The Debasing Demon." Analysis **70**(2): 228-237.
- Schiffer, S. (1996). "Contextualist Solutions to Scepticism." Proceedings of the Aristotelian Society **96**: 317-333.
- Searle, J. R. (1983). Intentionality: An Essay in the Philosophy of Mind, Cambridge University Press.
- Segal, G. (2000). A Slim Book about Narrow Content. Cambridge, MIT Press.
- Sellars, W. (1963). "Abstract Entities." The Review of Metaphysics **16**(4): 627-671.
- Smart, J. J. C. (1973). An Outline of a System of Utilitarian Ethics. Utilitarianism: For and Against. J. J. C. Smart and B. Williams. New York, Cambridge University Press.
- Sosa, E. (1999). "How to Defeat Opposition to Moore." Noûs **33**: 141-153.

- Sosa, E. (2000). "Skepticism and Contextualism." Philosophical Issues **10**(1): 1-18.
- Stanley, J. (2005). Knowledge and Practical Interests. New York, Oxford University Press.
- Stanley, J. and T. Williamson (2001). "Knowing How." Journal of Philosophy **98**(8): 411-444.
- Stoutenburg, G. (2015a). "The Epistemic Analysis of Luck." Episteme **12**(03): 319-334.
- Stoutenburg, G. (2015b). "Think of the Children! Epistemic Justification and Cognitively Unsophisticated Subjects." Pacific Philosophical Quarterly.
- Stoutenburg, G. (2015c). "Vicious Regresses, Conceptual Analysis, and Strong Awareness Internalism." Ratio.
- Tucker, C. (2016). Acquaintance and Fallible Non-Inferential Justification. Intellectual Assurance: Essays on Traditional Epistemic Internalism. B. Coppenger and M. Bergmann. New York, Oxford University Press.
- Unger, P. (1971). "A Defense of Skepticism." The Philosophical Review **80**(2): 198-219.
- Unger, P. (1975). Ignorance: A Case for Scepticism. New York, Oxford University Press.
- Williamson, T. (2000). Knowledge and its Limits. New York, Oxford University Press.
- Williamson, T. (2005). "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge." The Philosophical Quarterly **55**(219): 213-235.
- Wright, C. (1983a). "Keeping Track of Nozick." Analysis **43**(3): 134-140.
- Wright, J. P. (1983b). The Sceptical Realism of David Hume. Minneapolis, University of Minnesota Press.