INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

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

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

SKEPTICISM AND EPISTEMIC AGENCY

by
Jill Rusin

A dissertation submitted to The Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
April, 2002

UMI Number: 3046552

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Epistemic contextualists like David Lewis allow that we have substantially *infallibilist* reflective intuitions about knowledge even though our everyday talk accepts *fallibilist* attributions of knowledge. They give serious weight to both our everyday talk and our propensity to assent to the skeptic's conclusions, and give us a concept of knowledge that accommodates both. The skeptic would, of course, leverage such infallibilist intuitions in order to undermine the legitimacy of our everyday attributions. Most contemporary epistemologists would simply argue that our concept of knowledge is fallibilist, but this doesn't give them a particularly convincing reply to intuitively compelling skeptical arguments.

This essay, most generally, takes up the question of how best to argue about skepticism. I claim that a philosophical account of knowledge that responds to skepticism should be able to clearly accommodate our epistemic agency. In chapter 1 I explain how skepticism challenges our naïve conception of agency and introduce two criteria for an adequate response to skepticism. In chapter 2, I discuss contextualist solutions; I argue they fail to satisfy our criteria *and* fail on their own terms; I explain how both failures relate to a failure to adequately treat epistemic agency. In chapter 3 I develop a conception of epistemic agency that describes what it is to be an active, reasoning source of belief. This picture is at odds with a dominant strain of epistemology that has it that one is substantially subject *to* belief, not effectively responsible for it; I criticize an influential argument by Bernard Williams that claims belief is not an attitude we control. This discussion in chapter 3 gives support to, a rationale for, the anti-insulationism requirement set out in chapter 1. Finally, in chapter 4 I return to the

methodological issues of how to argue about skepticism. I first defend the idea that a knowledge-based skepticism is worth philosophical attention. I then offer an interpretation of what the skeptic is up to, and discuss what an effective counterstrategy might look like; I suggest that a particular sort of *normative* argument is the best way with the skeptic.

Committee: Professor Michael Williams (philosophy)
Professor Richard Bett (philosophy)
Professor Richard Flathman (political science)
Professor John Marshall (history)
Professor Ali Khan (economics)

Table of Contents

Introduction

To spend time thinking about skepticism has always invited ridicule. And not just from the man on the street. In the preface to his recent book about skeptical arguments, John Greco says: "This book is largely the result of snide remarks from my colleagues at Fordham University. I arrived there very interested in skeptical problems, and people...would wonder, out loud, why."[1] I share Greco's interest in skeptical problems. Nonetheless, I, too, might ask him why (though likely enough with a different tone). This seems like a perfectly good question to ask. I'm not sure I would be quite as interested in skeptical arguments if I were inclined to think of them in the way that Greco, along with many other epistemologists, thinks of them.

The tendency (at least since Hume) is to consider skepticism a wholly theoretical problem. It is a problem in theory, but not in practice; nobody professes to *be* a (philosophical) skeptic these days. And I don't, either. But I do think that skepticism must bear some more interesting connection to the everyday that's worth investigating. I find the skeptic's puzzles are so difficult, and so interesting, for the same reason: the skeptic offers normatively appealing epistemic principles. And the appeal of these principles can be practical as well as theoretical. If as theorists, we work hard to find a way past them, as participants in everyday epistemic life we can find ourselves under their sway.

So I argue. In doing so, I take up a type of skeptical problem which to many epistemologists, even some of those most interested in skepticism, is no problem at all. Stuart Cohen has described this problem as a paradox that makes trouble for our ability to

[1] John Greco, *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry* (Cambridge: Cambridge University Press, 2000), xiii.

be contented fallibilists. It is now best known as the "Argument from Ignorance," formulated as follows:

The Argument From Ignorance (AI)[2]

1. I don't know that not-H.

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

C. I don't know that O.

H is some skeptical hypothesis, incompatible with *O*; *O* is some ordinary fact, ordinarily thought to be known. In this form, it has been at the center of discussion of contextualist approaches to skepticism, and contextualist analysis of our knowledge-talk. These contextualist views themselves have been the subject of a significant (and increasing) amount of discussion within epistemology in the last ten years or so. But, although the contextualists discuss a *form* of skeptical argument I find interesting, I cannot endorse their solutions. At the core of my dissatisfaction with contextualist views are two broader issues: (1) the problems for *epistemic agency* that contextualism leaves us with, (2) the *methodological* approach to skepticism the contextualists take.

The contextualist methodology, in brief, is as follows. The contextualists don't take skepticism to actually pose some problem of the *validity* of our everyday knowledge attributions. This is how the skeptic wants to be understood, and what it might be to take skepticism quite seriously. The contextualists, though, take our everyday attributions to be all right, and take skepticism to pose a puzzle for the *theoretical* project of *understanding* our knowledge-talk. This is a descriptive project of analysis, and our everyday talk is data for this project.

---

[2] This formulation comes from Keith DeRose, "Solving the Skeptical Problem," *Philosophical Review* 104 (1995): 1.

It is the case that our everyday talk accepts *fallibilist* attributions of knowledge. That is, we accept that a person can know for a reason, or on some basis, which does not *entail* the truth of the belief it supports. Most all of the time when we attribute knowledge to some subject *S* that some proposition *p*, it is the case that *S*'s basis for belief in *p* is defeasible. But our susceptibility to skeptical argument is also data for the contextualist. He takes seriously our tendency, when faced with a skeptical argument like AI above, to arrive at the skeptic's conclusion; he takes seriously what can be characterized as substantively *infallibilist* reflective intuitions about knowledge. We are generally not pleased to accept open-eyed attributions of knowledge to S *where* and *when* it is pointed out to us the not-*p* possibilities *S* has failed ruled out. The contextualist allows the truth conditions for knowledge attributions to be governed by both our everyday fallibilism and our more infallibilist reflective intuitions.

I sketch, broadly speaking, these features of the contextualist view here, in order to show how the various concerns of this essay are connected. Most generally, this essay takes up the question of how to think about skepticism—how to conceive of it, how to approach it, how best to argue about it. I argue for specific answers to these questions, but it also my intent in pursuing these methodological questions to contribute to our understanding of the relation between our everyday beliefs about knowledge and our philosophical concepts of it. Alongside these more methodological questions is an argument for the importance of considering epistemic agency in connection with skeptical argument. Our everyday conception of our epistemic agency bears important connections to our everyday concept of knowledge; I argue that a philosophical account of knowledge that responds to skepticism should be able to clearly accommodate our

agency. Considering contextualism provides an opportunity to make these connections, and gives a motivation to pursue both the questions about agency and the questions about methodology. For the contextualist ends up with a view that is problematic for epistemic agency, and I'll argue that the best way to avoid this problem is by taking a different methodological approach. Looking at the contextualist view shows us why more attention is needed to these two areas, if we want a satisfactory response to skepticism, if we want a better understanding of it.

The structure of the essay is as follows. In chapter 1 I explain how skepticism challenges our naïve conception of our epistemic agency. The skeptic invites us to ask: (1) what norms underlie our concept of knowledge, or applications of our concept of knowledge; (2) are these norms we can (reflectively) endorse? and also (3) are we responsive to these norms, capable of responsibility to them? I introduce a preliminary sketch of our ordinary conception of epistemic agency, and introduce two criteria for an adequate response to skepticism when viewed as a threat to epistemic agency. One is an anti-insulationism requirement: an adequate theory will allow sufficient integration of our practical and reflective agency. It should be possible to see that our reflective selves can endorse the epistemic norms we are governed by. The second is that an adequate account will be an account of *our* epistemic concepts, that is, epistemic concepts grounded in values we endorse. Next, in chapter 2, I discuss contextualist solutions to skepticism. I argue that such accounts clearly fall afoul of our first criterion on a satisfactory response to skepticism, and also that there is reason to doubt that they satisfy the second criterion. I explain how these failures both relate to a failure to adequately treat epistemic agency, a topic I go on to explore more fully in chapter 3. Chapter 3 expands on the conception of

agency introduced in chapter 1. Drawing from the ideas of Philip Pettit and Michael Smith's *conversational stance*, Harry Frankfurt's concept of *identification* and Richard Moran's work on the distinctiveness of the *deliberative* outlook to first-personal rationality, I develop a conception of epistemic agency that describes what it is to be an active, reasoning source of belief. This picture is at odds with a dominant strain of epistemology that has it that one is substantially subject *to* belief, not effectively responsible for it; I criticize an influential argument by Bernard Williams that claims belief is not an attitude we control. This discussion in chapter 3 gives support to, a rationale for, the anti-insulationism requirement set out in chapter 1.

Finally, in chapter 4 I return to the methodological issues of how to argue about skepticism. I first defend the idea that a knowledge-based skepticism is worth philosophical attention against the arguments of Hilary Kornblith and Michael Williams, who argue that this kind of skepticism is uninteresting. I then offer an interpretation of what the skeptic is up to, and discuss what an effective counterstrategy might look like. I think that David Lewis (a contextualist) is right to note the disparity between our eyes-closed fallibilism and our eyes-open infallibilism about knowledge; we do have substantially mixed fallibilist and infallibilist intuitions about knowledge. But since I've argued that the contextualists' bi-perspectival solution is incompatible with a conception of agency we'd do best to retain, I need a different accommodation of Lewis's observation. I suggest that epistemological debate is unnecessarily encumbered by its methodological self-conception. The debate over skepticism (or more generally, over an account of our concept of knowledge) shouldn't be considered merely a project of adequately analyzing standing epistemic concepts. I take up a suggestion of Sally

Haslanger's that we instead pursue a project she calls critical immanent epistemology; such a project sees descriptive analysis best developed in tandem with questions about what epistemic concepts or properties we should value. I don't develop and defend answers to these valuational questions here, but given our discussion, the value of a fallibilist concept of knowledge is clear enough. We don't have to argue, then, that our concept of knowledge simply *is* fallibilist, or does not open itself to reflective infallibilism. We can argue for reflectively fallibilist intuitions not by saying we always already have them, but by arguing that we should. After all, the skeptic is most fundamentally arguing that we should *not* have fallibilist intuitions when it is the case that we do. The way to meet a normative argument is with a normative argument; this is the best way against the skeptic.

Chapter 1

The Skeptic's Assault on Epistemic Agency

> *'Are you in a position to know?': that is, you must undertake to show not merely that you are sure of it, but that it is within your cognizance.* — J.L. Austin

In this chapter, I first discuss the current tendency to view skepticism's significance as methodological. I argue that this is apt to leave underappreciated the normative significance of the skeptic's challenge. Where the skeptic's challenge is considered an actual normative problem, the threat is to our epistemic agency. I discuss what this means, what is at stake in considering the skeptic's challenge in this way, and discuss the criteria for an adequate response to skepticism, once it is explicitly considered as a normative problem.

1.1 The point of addressing skepticism: from Stroud's plea for significance to current claims of 'methodological import'

Barry Stroud's 1984 book on skepticism made a case for what he called skepticism's *significance*. At a time when skepticism was paid relatively little attention, Stroud argued that a comfortable solution to the skeptic's challenge was not at hand; the complacency, dismissiveness, boredom or impatience shown the topic by most philosophers, and many epistemologists, was misplaced, premature at best. Nearing twenty years later, thanks in part to the influence of Stroud's book, skepticism has regained a more prominent place within epistemology. What is notable about this shift is not that skepticism is perhaps now more discussed, more a topic for analysis; certainly there were many, many analyses and would-be refutations out there by 1984. Rather, Stroud's point about skepticism's significance seems to have been absorbed; there is a

correspondingly changed view of the explanatory burden borne by those who address skepticism. An attitude shift has taken root.

What is most interesting about Stroud's book is that he argues for more respect and attention for skepticism itself. Rather than champion a favorite solution, Stroud champions the philosophical problem, the significance of its difficulty. And while he does often argue on the skeptic's behalf, against various supposed solutions or dissolutions, he does not exactly argue *for* skepticism, either. That is, he does not straightforwardly argue for skepticism's *truth*. Stroud discusses historically prominent and theoretically representative approaches to disarming the skeptic, carefully arguing how and why each comes up short. But Stroud's book doesn't just seek to show that each attempt to refute skepticism fails, and it doesn't rush to offer any other neat solution. It urges some taking stock, the better to learn from all these heady attempts gone wrong. This strategy helped turn philosophical attention to a more cautious stance, towards reflection on what sort of problem skepticism is, what a satisfactory response to skepticism could even possibly be, what *sort* of answer we should look for.

Today, relatively few philosophers are willing to claim an easy victory over skepticism, or to simply argue we have nothing to learn from further engaging with it. Enthusiasm has faded for 'dissolving' the problem as meaningless, in the manner of the logical positivists or ordinary language philosophers. Instead there is recognition that skepticism is a complex problem, and that no simple solution can be expected to lay it to rest. Even those theorists who offer what is perhaps currently the closest thing going to a neat 'solution', the epistemological contextualists, admit that in addition to explaining their favored solution, in order for it to succeed, they also owe an explanation of *why*

there seemed to be a problem about knowledge in the first place. Otherwise, their solution gets little grip, they are well aware. Here is Stewart Cohen:

> What we want is a resolution of the [skeptical] paradox that preserves our strong intuition that we know things. But any such resolution must explain the undeniable appeal of skeptical arguments...because the paradox arises within our own thinking about knowledge—the premises of skeptical argument are premises we are inclined to accept—any successful response to the paradox must explain how we end up in this situation.[1]

And contextualists, with other contemporary epistemologists, realize that this is no small feat. A 'solver' of skeptical paradox in the traditional vein, albeit aware enough of this explanatory burden, the contextualist is offering what Steven Schiffer has called a "happy-face solution" to skepticism. Such a solution resolves skeptical paradox by rejecting one of the seemingly plausible propositions that give rise to it. As Schiffer remarks: "the typical trouble with happy-face solutions, a trouble well illustrated by the compatibilist's solution to the problem of free will, is that it leaves one wondering, 'if that's the solution, then what the hell was the *problem?* ' The happy-face solution makes it a mystery why one was ever deceived...in the first place."[2] Both Cohen and Schiffer's remarks recognize that the *prima facie* intuitiveness and intelligibility of the skeptic's premises make his argument difficult to combat. These last twenty years have witnessed a growing appreciation of this point; in some instances, this recognition has issued in changed expectations for a resolution.

---

[1] Stewart Cohen, "Contextualism, Skepticism and the Structure of Reasons," *Philosophical Perspectives* 13 (1999): 63; see also Keith DeRose, "Solving the Skeptical Problem," *The Philosophical Review* 104 (January 1995): "In seeking a solution to this puzzle, we should seek an explanation of how we fell into this skeptical trap in the first place, and not settle for making a simple choice among three distasteful ways out of the trap. We must explain how two premises that together yield a conclusion we find so incredible can themselves seem so plausible to us. Only with such an explanation in place can we proceed with confidence and with understanding to free ourselves from the trap" (3).

[2] Stephen Schiffer, "Contextualist Solutions to Scepticism," *Proceedings of the Aristotelian Society* 96: 329.

A lot of good work has gone into rooting out problematic assumptions embedded in particular formulations of skeptical argument, but it is harder to show that all skeptical arguments depend on or grow from these assumptions. Debunking Cartesian or Humean philosophy of mind or metaphysics, for example, does not necessarily defuse skepticism. Transcendental arguments, like those of Davidson or Putnam, have added important results to the discussion, but are not offered, by their authors, as 'solutions' to skepticism. Philosophy's approach to skepticism, and probably epistemology in general, has grown more sophisticated; a more democratic spirit prevails. While naturalized epistemology, for instance, was at one time seen by its proponents as doing away with problems like skepticism, at least as traditionally conceived, I am not so sure that its more contemporary adherents take the view that their research program entails the fruitlessness of alternative approaches to epistemology. If skepticism is increasingly seen as a multi-faceted, complex problem, then progress can be made in many different ways on different fronts. So while it is hard to say that there is any one favored approach to skepticism currently predominant, this is all to the good. A catholic attitude, along with a prevailing note of caution, is a real sign of philosophical progress.

So it is here, rather than in his more particular views (about skepticism's conditional correctness, or suggestions of skepticism's source in the epistemologist's project itself), that Stroud's book has had the most influence. Epistemologists now understand the argumentative and explanatory burden of those who address skepticism, and also understand that what stands to be gained is not so much vindication of "knowledge", but a deeper understanding of some or all of: our concept of knowledge, objectivity, epistemic agency, epistemology, our ordinary and philosophical assumptions. I think

many epistemologists would now endorse Michael Williams's aim of offering what he calls a 'theoretical diagnosis' of skepticism.[3] Theoretical diagnosis maintains the positive goal of turning away the skeptic's conclusions, and aims to do this by showing the skeptic to embrace theoretically optional or contentious views. But Williams too realizes the intuitive pull of the skeptic's argument, an argument intuitive enough that previous attempts to defuse it have been unsatisfactory. So, a successful theoretical diagnosis will need to be, indeed, a "*deep* diagnosis of its sources."[4] The contemporary theoretical diagnosis is not any simple solution. Given that even those theorists who ultimately hope to prove skepticism unfounded recognize that its sources run deep, then, there seems the promise of much to be learned from an investigation of skepticism. For either there is a truth in skepticism, stemming from something in our concepts or ourselves, or its deep source is instead to be found in philosophical assumptions or theorizing, making skepticism, as a problem, an artifact. But even in this case, because of the depth of our engagement with the problematic theory, what we stand to learn about ourselves remains significant.

So far I have documented the fact that there has been a notable shift in attitude towards skepticism. Contemporary epistemologists are less likely to embrace simplistic dissolutions. They recognize that insofar as a brash pragmatism or an unrepentant naturalism fails to address skepticism so much as do away with it, it cannot be seen as an adequate response to skepticism, as having *already* moved successfully beyond it. The

---

[3] See Michael Williams, preface to *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism* (Princeton: Princeton University Press, 1996), xii-xxiii.
[4] Williams, *Unnatural Doubts*, xvi. My italics. See also xix: "this will involve going very deeply into the theoretical background of skeptical problems...theoretical diagnosis is not to be equated with hasty dismissal."

prima facie intuitiveness of the skeptic's arguments engenders an explanatory burden. On the positive side, perhaps there is more to be gained via the work of such explanatory accounts. The contemporary epistemological climate, then, reflects more widespread acknowledgement of the significance for which Stroud argued.

One oft-repeated contemporary characterization paints this significance as "methodological". In their recent anthology, Sven Bernecker and Fred Dretske introduce a section on skepticism with the commentary: "Almost nobody thinks that skepticism is true. That does not mean skepticism is not important. Its relevance is methodological."[5] John Greco makes the case for skepticism's methodological import a major theme of his recent book:

> I argue that the analysis of skeptical arguments is philosophically useful and important. This is not because skepticism might be true and we need to assure ourselves that we know what we think we know. Neither is it because we need to persuade some other poor soul out of her skepticism. Rather, skeptical arguments are useful and important because they drive progress in philosophy.[6]

On one reading of such remarks, they are completely consonant with the points I have been making thus far. Saying that skepticism is methodologically important, these theorists endorse the significance of skepticism, and acknowledge the potential profits of engaging with it. Neither author considers the *truth* of skepticism to be at issue, but even though they do not think skepticism could be true, still they champion its role in epistemology. They agree we stand to learn about our epistemic concepts, or about our theoretical assumptions, or both. I do, however, want to look more closely at what is ruled out in this casting of skepticism's significance as methodological.

---

[5] Sven Bernecker and Fred Dretske, "Scepticism: Introduction," in *Knowledge: Readings in Contemporary Epistemology*, ed. Sven Bernecker and Fred Dretske (Oxford: Oxford University Press, 2000), 301.
[6] John Greco, *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and their Role in Philosophical Inquiry* (Cambridge: Cambridge University Press, 2000), 2-3.

What these authors rule out is open-ended inquiry, or live worry, concerning the truth-status of the skeptic's startling conclusions; whether skepticism might actually be true is not at issue. The challenge is rather to their own epistemological enterprise, to epistemological theories. It is to *show* that skepticism is false, and to show that epistemological theories under consideration are not committed to skepticism. What seems reasonable here is that, while Stroud may have made the case for the intuitiveness of skepticism, it is equally true that no one begins, or is simply, naturally, a skeptic. We *do* ordinarily think that we know a great many things. This much is fact. And it is also fact that it is incredible to think that the skeptic's dire conclusions could actually obtain. Although defenders of skepticism may urge we not 'beg the question' against the skeptic, it is hard to imagine that we could, or should, start from a completely neutral stance. It is no surprise that an epistemologists' goal would be to show *that* skepticism is wrong, not to investigate *whether* it is wrong. As human beings, we are, all of us, it seems, already committed to skepticism's being wrong.

But rather than pursue this line of reasoning, I want instead to go on to question the characterization of skepticism's significance as methodological. I am less concerned with the specific views of Greco, or Bernecker and Dretske, or any other particular theorist, more interested in the inclination towards this kind of characterization, and the implications stemming from it. I will argue that there is still much to be objected to, when, as epistemologists, we begin by considering skepticism's importance to be *merely* methodological. Doing so: a) rules out that there is *any* truth in skepticism and b) cuts short the possibility of normative effect or import flowing from discovery or recognition of this truth, or alternatively from a failure to come up with a satisfactory explanation of

why skepticism is *not* true. a) and b) combined seem to preempt the possibility of any very revisionary outcome to an encounter with skepticism. Rather than *start* with such a circumscribed outlook, I want to argue that skepticism is paradigmatically the type of problem that seems to demand normative attention. It poses a challenge to our rational selves, to the idea of our normative governance. It forces difficult normative questions: what should we believe, and on what basis?; when should we claim to know?; what is responsible epistemic behavior, and is there is such a thing?; what is important or valuable about our epistemic practices?[7] That skepticism is centrally a normative problem may seem obvious; the skeptic hopes that it is.[8] But if skepticism is *not* considered in *any* way a live option, if there can be nothing like a truth in skepticism, if it is instead approached as merely a theoretical exercise to show how it is that skeptical arguments are in error, then skepticism never gets its teeth as a normative problem. The merely methodological view (MMV) proves hasty. By circumscribing the characterization of skepticism's significance from the start, it limits the possible outcome of philosophical investigation. It also obscures the normative dimension of skepticism. And even the instrumental ends the MMV advocates (furthering philosophical progress) will be better served by bringing out, rather than obscuring, the normative aspect of the skeptic's challenge—much stands to be gained, philosophically, by considering skepticism specifically in terms of the normative threat it potentially poses.

---

[7] The first questions may seem more obviously provoked than the last. I hope to show how all these questions grow from skeptical challenges. For an excellent discussion of the relevance of this last question for epistemologists investigating knowledge, see Sally Haslanger, "What Knowledge Is and Ought to Be," *Philosophical Perspectives* 13 (1999): 458-480.

[8] Well, in one sense, and depending on what 'skeptic' one is talking about: you might say the skeptic doesn't hope, because he doesn't care. Here I only meant to link seeing the problem as obviously normative with a holding open of the thought that skeptic has something important to say--that it isn't decided to begin with, or regarded as manifest, that skepticism presents a neat puzzle for epistemologists and nothing deeper.

I now want to elaborate on a) and b), first looking at what finding 'some' truth in skepticism might come to, by surveying what some more skeptically-inclined theorists have had to say here, and then discussing what manner of normative challenge skepticism potentially presents.

1.2 Truth in skepticism?

At least on the face of it, the merely methodological view is at odds with the conclusions of a number of philosophers who have worked on skepticism. These philosophers have conceded some ground to the skeptic, agreeing with him at least in part. According to these theorists, all our pretheoretical ideas about knowledge cannot be jointly, coherently sustained. The skeptic is correct about this much. MMV, however, seems to rule, in advance, that these philosophers are not right.

Steven Schiffer says: "our concept of knowledge has a deep-seated incoherence. The criteria we employ to tell us when we know something conflict, and the concept contains no higher criterion whose application can resolve the conflict."[9] He discusses the skeptical paradox formed by an observation instance of ordinary Moorean fact, an instance of the closure principle, and observation of an uneliminated skeptical hypothesis:

1. I know that I have hands.
2. If I know that I have hands, then I know that I'm not a handless-brain-in-a-vat (BIV).
3. I don't know that I'm not a BIV.

Of course, 1 and 2 yield 3*: *I know I'm not a BIV.* 1, 2 and 3 are mutually inconsistent. But Schiffer says that all three embody criteria we use in ascribing knowledge, that none can simply be given up. We count as knowing things like 1) because, roughly, we count

[9] Schiffer, "Contextualist Solutions to Scepticism," 330.

as knowing beliefs delivered by our senses when they are well-functioning, when the beliefs are true, and when there is no special reason to doubt them. We are entitled to 2) because we know what is deduced from premises we know to be true. 3 holds because: "we don't count ourselves as *knowing* that the defeating etiology [BIVdom] doesn't obtain unless we have evidence for its not obtaining that goes beyond what one normally has when one takes oneself to have perceptual knowledge."[10] Again, here is Schiffer's pessimistic conclusion:

> This is...the incoherence in our concept of knowledge. No wonder we're of two minds about the sceptic's claim not to know that she's a BIV, and no wonder deep analysis of our concept of knowledge provides no release for this ambivalence. It's why the sceptical paradox has no happy-face solution.[11]

Thomas Nagel finds the incoherence to be between what he calls 'objective' and 'subjective' viewpoints that cannot be jointly sustained. He takes skeptical possibilities seriously, and thinks that we will not be able to show that skeptical worries are incoherent worries. Objectively, skeptical possibilities are real and uneliminated and our beliefs outstrip their support. But subjectively, we cannot sustain this view.

> The objective standpoint here produces a split in the self which will not go away, and we either alternate between views or develop a form of double vision...when we view ourselves from outside, a naturalistic picture of how we work seems unavoidable. It is clear that our beliefs arise form certain dispositions and experiences which, so far as we know, don't guarantee their truth and are compatible with radical error. The trouble is that we can't fully take on the skepticism that this entails, because we can't cure our appetite for belief, and we can't take on this attitude toward our own beliefs while we're having them...there is no way of bracketing our ordinary beliefs about the world so that they dovetail neatly with the possibility of skepticism.[12]

Nagel's view is skeptical. He thinks that the subjective and objective viewpoints cannot be combined, and neither could be given up. Moreover, this is no contented 'viewpoint pluralism'—there is the suggestion that although the subjective point of view is real, it is

---

[10] Ibid., 332.
[11] Ibid., 332-3.
[12] Thomas Nagel, *The View From Nowhere* (New York: Oxford University Press, 1986), 88.

the objective view that gets things right, that normatively, the objective view of things is in the right. If this is the case, there is a genuine sense in which our pretensions to knowledge outstrip our abilities to know; the skeptic is vindicated.

Barry Stroud argues, beyond the defense of skepticism's significance we have discussed, in addition for skepticism's 'conditional correctness': "the skeptical conclusion would be correct if the philosophical question to which it is an answer were legitimately posed."[13] Of course, 'conditional correctness' is a cautious position vis-à-vis skepticism. It does not admit that the skeptic's question *is* legitimately posed, or fully natural. And indeed, Stroud's writings since *The Significance of Philosophical Scepticism* more strongly suggest that it is the epistemological project itself that is the source of skepticism. But even here, Stroud is never explicit that this very project is illegitimate. If there is something peculiar about the desire to understand knowledge in general which drives towards skepticism, it remains to show that this desire *cannot* be fully coherent, or fully natural. Stroud insists enough on the seeming meaningfulness of the skeptic's questions against skepticism's potential dissolvers that he still gets provisionally classed with those who hold there is some truth to skepticism.

Peter Unger has argued that skepticism is true, Stanley Cavell has said that there is "a truth in skepticism", Robert Fogelin defends a Pyrrhonian skepticism.[14] David Lewis, with other contextualists, says that within the epistemological context, skepticism is true[15]. For all these theorists, skepticism captures something of significance beyond what

---

[13] Barry Stroud, *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press, 1984), 256.
[14] Peter Unger, *Ignorance: A Case for Scepticism* (Oxford: Oxford University Press,1975); Stanley Cavell, *The Claim of Reason: Wittgenstein, Scepticism, Morality and Tragedy* (Oxford: Oxford University Press, 1979); Robert Fogelin, *Pyrrhonian Reflections on Knowledge and Justification* (New York: Oxford University Press, 1994).
[15] David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (December 1996). And additionally, for Lewis, there is also a sort of truth-in-skepticism that transcends the skeptical context:

it makes sense to call a merely methodological importance. And if skepticism is true, or if there is 'some' truth in skepticism along any of the various routes these philosophers take in making out what this might be, then it is not antecedently clear that this could be of importance *only* for philosophical theorizing. So it seems premature to say, with philosophers like Bernecker, Dretske and Greco, that we do *not* need to reassure ourselves that skepticism is not true, that no one believes that skepticism is true, or that its real importance can only be instrumental. At least not as long as we take the views of the truth-in-skepticism philosophers seriously.

1.3 The queerness of the merely methodological view

Perhaps it may be said, in defense of MMV, that we are overreacting to its characterization of skepticism; read charitably, perhaps it is rather an innocuous claim. It says *something* about the nature of skepticism's significance. That is, beyond siding with Stroud against the dismissers, it construes this significance, calling it 'methodological'. Perhaps simply, and most basically, this is to register only very pedestrian fact: if skepticism's importance is for philosophical theorizing, this is because people do not, or cannot, think skepticism is true, they cannot be motivated to act by consideration of skeptical arguments.

But this itself needs clarification. If the claim is only that *ordinarily* people do not believe skeptical claims—even the skeptic admits *this* verity. This does not conflict with the possibility of skepticism's truth, of the ability to become reflectively convinced of skepticism's truth, or the possibility of one's actions and attitudes being affected by

---

"Never—well, hardly ever—does our knowledge rest entirely on elimination and not at all on ignoring. So hardly ever is it quite as good as we might wish. To that extent, the lesson of skepticism is right—and right permanently, not just in the temporary and special context of epistemology" (563).

consideration of skeptical argument. While unobjectionable, this conservative rendering does not go very far towards *any* sort of indication of skepticism's significance. MMV certainly suggests somewhat more. Characterization of skepticism's significance as methodological, as driving progress in philosophy, does insinuate a contrast—*this* sort of instrumental significance with significance stemming from its being true, or *philosophical* significance with some sort of *practical* consequence beyond the theoretical realm. And neither of *these* contrasts follows simply from pedestrian fact. MMV, suggesting such contrast, must additionally assume either or both 1) we cannot be even theoretically, reflectively committed to skepticism's truth; 2) skepticism cannot have any practical impact.

There are arguments for 1), but it is rather a strange assumption for MMV to make; construed this way, MMV is odder than it first sounds. For, as we have noted, 1) excludes the views of truth-in-skepticism theorists. MMV says that skepticism is important because of what we can learn from it, from the manner in which it drives philosophical progress. And there is at present no broad consensus, no widely accepted diagnosis of skepticism agreed upon by philosophers. But certainly a lot of important work, some of the most subtle exploration and discussion of skepticism has been carried out by the truth-in-skepticism theorists. But if MMV assumes 1), it is committed to opposing the conclusions of these theorists.[16] Of course there is nothing wrong with the thought that we can learn from views that are false—it would not be incoherent for the MMV to say that the truth-in-skepticism views have taught us a lot, but are false. The

[16] In the case of some of these theorists, compatibility can perhaps be finessed—but the point is that to take the views of these theorists seriously is to reject the sort of characterization that MMV puts forward. Where 1) is not actually incompatible, because of the subtleties of explanation needed to show that they *are* compatible, it is at very least misleading as a characterization that could stem from these views. If we *start* with this characterization in mind, it is hard to see how we would *arrive* at this type of view.

problem is, rather, that if the truth-in-skepticism theorists have contributed to philosophical knowledge, they have done so by propounding views that the MMV appears to rule out from the start. If *these theorists* had accepted MMV, construed in this way, they would not have made the contributions they did, contributions that have served to advance philosophical theorizing. So at least some significant portion of the instrumental importance MMV ascribes to skepticism could not be achieved except insofar as certain philosophers are misconceived about skepticism's real import (the 'real import' per MMV). This makes MMV a strange view. MMV undercuts itself if it suggests that skepticism is significant because it furthers philosophical knowledge, yet precludes the approaches of those whose exploration of skepticism help contribute to that knowledge. Embracing 1) makes MMV a (partly) self-effacing meta-epistemological theory.

Another way of explaining the queerness of wedding MMV to 1) is to look at the arguments in support of 1). Such arguments can be divided into two camps—the naïve arguments, and the sophisticated arguments. The trouble is that the naïve arguments are too naïve to support the contention that skepticism *is* significant. And the sophisticated arguments are subtle enough, and require argumentative and explanatory support enough, that they hardly succeed in advance of carrying out of the philosophical investigation which MMV is supposed to characterize from the outset. Requiring this degree of argumentative support, MMV doesn't just begin to look a bit less like an innocuous characterization, more like a substantive viewpoint, but indeed becomes wedded to the whole of an analysis that would bear it out, make good on its assumptive claims. So we could accept it if we *already* accepted the whole of a particular analysis of skepticism

compatible with it; otherwise, it is unnecessarily constricting—there is little reason to swallow it, and, I argue, positive reason to avoid it.

The naïve argument for 1) is the reaction to skepticism that Keith DeRose calls the "Aw, Come On!" response. It is not so much an argument that we cannot be reflectively convinced by skepticism, as a psychological reaction to skeptical argument, a feeling that one is not so convinced in the least. Here is DeRose on this camp:

> ...finding the arguments farfetched, ridiculously weak, and quite unthreatening...such a reaction is often accompanied by an exclamation somewhat along the lines of, "Aw, come on!" Those inclined to react in this latter way probably grew increasingly impatient of my repeated description...of the Argument by Skeptical Hypothesis as "powerful," thinking instead that the argument hasn't a chance in the world of establishing its absurd conclusion.[17]

But this view, or reaction, does not yet accede to Stroud's point, discussed above, that skeptical arguments are significant because they really do seem to be intuitive, powerful arguments. It is more dismissive than MMV.

There is a more sophisticated argument available in support of 1), which can accommodate Stroud's point, while still supporting the basic feeling of the "Aw, Come On" camp. This approach admits that the premises of skeptical argument do seem compelling, but buttresses the naïve "come-on!" intuition with a type of Moorean argument against the possibility of being even theoretically convinced by skeptical argument. David Lewis says that the fact that we know many things is "a Moorean fact...to doubt this in any serious and lasting way would be absurd...It is one of those things that we know better than we know the premises of any philosophical argument to the contrary."[18] But Lewis's strategy is actually more sophisticated than Moore's—he does not *just* argue, in Moorean fashion, that we are, and must reasonably remain, more

---

[17] Keith DeRose, "Responding to Skepticism," in *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted Warfield (New York: Oxford University Press, 1999), 3.
[18] Lewis, "Elusive Knowledge," 549.

confident of the fact that we know many things than of the conclusion to any sort of philosophical argument. Lewis, in line with the contemporary consensus I outlined above, recognizes the power of the skeptic's arguments. So he also recognizes the explanatory burden that follows on the heels of taking such a Moorean stance. And he provides an account to this end: "I started with a puzzle: how can it be, when his conclusion is so silly, that the sceptic's argument is so irresistible?...my [account was] built to explain how the sceptic manages to sway us."[19] Lewis's calling the skeptic's conclusion 'silly' is not just an 'Aw, come on!" response insofar as he takes on the burden of explaining why the skeptic's argument does not lead to a conclusion which robs us of our ordinary knowledge claims, showing them to be largely false. Lewis *is* committed to preserving most of our knowledge, most of our ordinary judgments about who has knowledge. This is because he thinks this is what there is *to* explain, in discussing skepticism.[20] To agree here, one needs to accept Lewis's account; the assumption of this optional point of departure is made good if one agrees with his analysis. Because *pace* Lewis, there *are* other options—one might say, with Schiffer, that we can't make coherent sense of all our pretheoretic intuitions about knowledge, that reflection reveals no way to make consistent all the conflicting rules underlying our evaluative, knowledge-ascribing practices. Even if one's aim is specifically to make sense of our epistemic practices and concepts (rather than to concede, skeptically, that

---

[19] Ibid., 561.

[20] See Lewis, "Elusive Knowledge," esp. pp. 560-561, 563, 549. Here I'm agreeing with Jim Joyce's description of Lewis's methodological point of departure: "The Lasting Lesson Of Skepticism," unpublished manuscript. Many others are in agreement with Lewis; here, for example, is Greco: "skeptical arguments are important not because they might show that we do not have knowledge, but because they drive us to a better understanding of the knowledge we do have" (*Putting Skeptics in Their Place*, 3). Greco operates from the same methodological point of departure that Lewis does.

this cannot be done), one needn't assume, with Lewis, that most of our ascriptions are true.

Peter Unger, in *Philosophical Relativity*, puts forth an 'invariantist' account of the truth conditions for knowledge attributions: to have 'knowledge', according to this account, means that you are in so strong an epistemic position with respect to the subject at hand that there is no better position to be had. This sets standards for knowledge high enough that most all of the attributions of knowledge we ordinarily make are false. Unger, who also gives an alternative contextualist account of 'knowledge', argues that the invariantist account has just as much in its favor as the contextualist theory; explanatory considerations decisively favor neither option. So according to Unger, it is not simply the case, as Lewis would have it, that the truth of our ordinary knowledge claims is the explanandum to be squared with the skeptic's premises; for Unger it is actually indeterminate whether our attributions are true or not, because there is nothing to determinately fix their semantics.

Jim Joyce also argues (and directly against Lewis) that there *is* choice when it comes to best 'saving the phenomena' of our pretheoretical epistemic talk and judgments. Unlike Unger, he doesn't argue for relativity and indeterminacy. In opposition to both Lewis and Unger, Joyce thinks that explanatory considerations favor a semantics and pragmatics that side with the skeptic's assessment of the falsity of our ordinary attributions. Rather than save the truth of most of our ordinary attributions and give up *other* commonsensical ideas about knowledge,[21] Joyce argues that we can consider such attributions substantially false, given a good explanation of their purpose. Joyce endorses

[21] And to be sure, Lewis does give up certain commonsensical intuitions about knowledge, more than he lets on. I will have more to say on this point below; see chapter 2, section 8.

Lewis's contextualist account of the truth conditions of knowledge attributions when Lewis's contextualist rules are recast as rules of warranted assertibility; he explains why 'hyperbolic' attributions of knowledge serve our purposes well—for example, they suggest when to terminate inquiry.[22]

The sophisticated argument, then, for 1) is an argument for a particular methodological point of departure, the one that Lewis takes up in the face of skeptical argument. It is an argument, really, for what the results of conceptual investigation *could* tell us, or could not, about knowledge.[23] If the skeptic's argument seems to land us in paradox, wedded to a conclusion we cannot swallow, Lewis, with Moore, is committed in advance to finding an explanation that rejects that conclusion. Lewis wants to square our epistemic practices and everyday usage with skeptical argument that seems to make a mess and a mockery of them. He assumes that an explanation that finds our attributions substantially false would not be an explanation of those practices. But I have tried to argue, briefly, that this is a more complicated and more controversial assumption than Lewis lets on. It assumes that those practices *can* successfully be sorted out, and perhaps also assumes that there is an associated concept of knowledge that can be sorted out. Given acknowledgement of the difficulty of sorting out skeptical paradox, it is not antecedently clear that this is the case. There is reason enough to be wary, to be concerned about the possibility that our pretheoretical intuitions about knowledge cannot

---

[22] Joyce cites Stroud's response to Austin as inspiring his approach in this manuscript; see Stroud, *Significance* p. 57 ff. on how conditions of appropriate assertion can come apart from truth conditions.

[23] So the transcendental arguments of Davidson and Putnam which suggest conceptual constraints on skeptical possibilities and conclusions could also be counted as arguments for 1). But I do not separately consider this type of argument, for my response would be parallel to what I say in the body of the text about Lewis's methodological approach, and in this case, I believe, my point is that much more obvious: if MMV is to be buttressed by this kind of argument, it becomes wedded to our acceptance of the whole of a larger account—then we do not have a working characterization of the significance of understanding or solving skepticism, but rather we are offered a solution itself.

be made coherent. Stephen Stich suggests (in the context of a broader discussion of philosophical method) that it is highly questionable to assume that we can use reflective equilibrium to sort out our epistemic concepts. He gives the following list of 'questionable presuppositions' of such a project: (a) there is *only one* such notion (e.g., of justification, or of knowledge) grounded in our practices; (b) our epistemic practices and the concept(s) they ground *are* each *coherent;* (c) a commonsense *concept* can be determinately separated from associated *'folk' theory* in some clear way; (d) some general *principle* or principles govern our concept (and our application of our concept), rather than some more variable, exemplar-based mechanism of application.[24]

I will have more to say about these matters below. Here, I want only to have shown that successful defense of 1) is dependent on the outcome of a larger analysis that makes good on its assumption and cannot stand apart from it. MMV, meant to innocuously characterize skepticism's significance, is not so innocuous when understood as holding that we could not be even reflectively convinced of skepticism. We should accept this only if we've *already* accepted the whole of a particular analysis of skepticism compatible with it; otherwise, we are not so obliged.

What about the second possibility for MMV—perhaps the intent of its characterization is rather to underline the fact that skepticism cannot have practical impact, that people cannot be motivated to act by consideration of skeptical arguments, even if they could be theoretically convinced of their merits. But here, too, there is

[24]Stephen P. Stich, "Reflective Equilibrium and Analytic Epistemology," in *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation* (Cambridge, MIT Press: 1990), 87-89. For Stich, the larger issue is why we should even *care* whether our beliefs are sanctioned by some evaluative epistemic concept that is 'ours'. Along the way to this bolder line of investigation, Stich gives us these reasons to doubt that we can arrive at a philosophical analysis of our epistemic concepts via reflective equilibrium.

something odd about this cashing out of skepticism's significance. Saying skepticism's importance is only theoretical and not practical would seem to imply not just the psychological claim that we are incapable of this kind of motivation, but also the descriptive and normative claim that this is not the sort of problem it is—not a practical problem. But assume that there is a metaphysical fact about whether skepticism is true or not. (Let us assume, for the moment, that here by skepticism's truth we mean, roughly, that we humans regularly fail to fulfill the conditions our concept of knowledge lays down for us.). This is separate from the question of what normative consequences flow from this fact. It is also separate from the question of whether we are capable of acting on or from the normative consequences that would flow from that fact, if true. Even if skepticism is *not* true, that does not mean that there may not still be reason to treat it as a normative, not wholly theoretical, problem—that is, a predicament from which normative consequences flow. To ignore the normative aspect of skepticism because we start with the conviction that the skeptic's conclusion is not true seems to mischaracterize the type of problem that skepticism is, or would be, if it is a real problem. If skepticism is a predicament from which normative consequences follow, then it is a normative problem (has a significant normative dimension) whether or not we could act on or be motivated or guided by those normative conclusions or consequences and whether or not the facts which would motivate the normative consequences actually do obtain.

The Merely Methodological View seems to have decided that 1) skepticism is false and that 2) skepticism is of no possible practical import. It is not clear whether 2) is supposed to follow directly from 1) or is an additional assumption. I have argued that MMV should not embrace 1). And now we have additional reason—it is too easy to

move from assuming that skepticism is not true, so wholly hypothetical, to thoughts that it is of no normative consequence, wholly theoretical, of no practical import.

But it is likely that MMV stems from the idea that, even if skepticism *is* true, we cannot believe it, so whether or not 1) holds, 2) should be uncontroversial. I have already argued that we need to be more careful in getting behind exactly what we mean to claim with such supposed platitudes as 'we cannot believe it', referring to the skeptic's conclusion. Again, if this is supposed to mean that ordinarily we do not believe it, the skeptic agrees here. But she thinks that her arguments, given due consideration, change things. At least some prominent voices have substantially agreed that there *is* no reflective escape from the skeptic's dire conclusions. So the platitude is recast as the thought that we cannot sustain this conviction, it cannot bear any weight for us, it cannot move or motivate.

But this, too, needs explanation if it is to support the idea that skepticism is a wholly theoretical problem. One thought is that 'practical import' is being construed in the following way: obviously, skepticism doesn't have practical importance, because even those who claim to believe it reflectively still act otherwise as normal. So how practically important is it, what discernible effect does it have?

Presumably 'acting as normal' means e.g. choosing to exit by the door, rather than the window or the ceiling, evincing ongoing commitment to first-order belief that skepticism impugns. I will have much more to say about this sort of dissatisfaction. For now, I want merely to point out that this kind of construal of 'practical effect' takes a not-very-subtle manifestation of effect as a paradigm. Of course there are subtler ways to have an effect. And one of those ways is a potential intellectual effect on our attitudes,

our self-concept, our self-regard, if the skeptic convinces us to view our knowledge claims as unwarranted. At stake is a view of our selves and our rationality, our rational self-respect. To characterize skepticism's significance as merely methodological seems to miss this out. If MMV is based on an assumption of 2), it suggests in advance that our reflective attitudes about the well-foundedness of our epistemic practices and the normative soundness of our ascriptions is necessarily and univocally motivationally inert. I think this is an unwarranted theoretical assumption.

If all the methodological view says is that we stand to learn a lot by engaging with skepticism, without prescribing in advance what that might be, and of what effect, then well and good. But then this is just saying that skepticism is significant. So why say this significance is specifically, wholly, methodological? Perhaps those who say this are not actually advancing the thesis that skepticism *itself* is important for this reason, but are really meaning to argue that skepticism's importance is the promise of one *particular* philosophical advance: his or her correct analysis. If this is the case, MMV is a corollary of the more particular view, and as a general claim mere pretense. Otherwise, we'd do best to drop the methodological characterization of skepticism's significance. Wedded to either the idea that skepticism cannot so much as be theoretically convincing, or the idea that it cannot have any practical import, it can only seem to preclude consideration of conflicting approaches and views. To do this is premature—a decision here seems to foreclose exploration of much of the philosophical terrain that promises to be of interest, and moves away from the catholic attitude I suggested (in the first section of this chapter) to be a real philosophical advance in thinking about skepticism.

Having described this danger inherent in the merely methodological characterization of skepticism's significance, I would now like to look more closely at the normative dimension it leaves underappreciated.

1.4 Skepticism as a normative challenge

If the skeptic's challenge really is profound, and not linguistic quibbling, it will raise questions about the propriety of our ordinary epistemic claims and practices. The skeptic offers what is supposed to be a shocking conclusion: we do not really know what we ordinarily think we know. If he is indeed serious about his skepticism and its shock value, his conclusion is not tethered to some rarified conception of knowledge, claiming only *that* ideal unattainable by human beings. His distressing conclusion does not place only infallible or indefeasible or certain knowledge beyond our reach, leaving us still with a lesser, workaday knowledge intact. Rather, his conclusions are supposed to be directly at odds with what we ordinarily think. So his target is more basic, more central to our everyday cognitive lives, germane to our *ordinary* self-conception; he is not just aimed at a particular *philosophical interpretation* of our cognitive lives or our knowledge. In just the sense that we mean it when we say of our self or of another that we know, the skeptic says we do *not* know. Or, since it can be awkwardly self-defeating for the skeptic to offer this positive conclusion, we can instead view the skeptic as presenting a paradox such that we cannot see *how* or *whether* we can know. This more careful version of skepticism nonetheless presents a serious challenge. Serious, not just because skepticism is a notoriously difficult and enduring problem of philosophy, but because what is at stake is, once the skeptic has thrown down the gauntlet, quite important to us, if not to him. Roughly, it is our own rational self-image, our epistemic

agency, epistemic normativity. To my mind, the challenge of skepticism is to see how we can maintain a view of our epistemic selves worthy of self-respect, despite what the skeptic has to say. If we cannot maintain our prior views of epistemic agency and normativity, the challenge becomes to see how we could go on, after they give way, to discover what self-image we are left with, and what might follow from it.

Chapter 3 of this dissertation will probe, in more detail, the picture of rational agency that is purportedly threatened by the skeptic's arguments. I will be particularly concerned to explore whether and how such a conception is committed to the idea that people have some sort of control over their beliefs, with arguments against the possibility of such control, and with the question of whether skeptical arguments themselves depend on objectionable assumptions of control. In the remainder of this chapter, I provide a preliminary sketch of the sort of rational agency threatened by the skeptic, and then discuss some constraints on a satisfying response to skepticism that emerge from looking at skepticism with this concern in the foreground.

1.5 Epistemic agents and epistemic agency

Humans are generally able to, and often do, have second-order attitudes toward their first-order beliefs. They also judge the beliefs of others. They can consider beliefs hasty, tenuous, dubitable, dubious, creditable, considered, reprehensible, appalling, well-grounded, shameful, shocking, reproachable, unassailable, unimpeachable—of course, the list goes on and on. This repertoire of attitudes, and the normative practices they help constitute distinguish human epistemic agents from epistemic wantons. By 'epistemic wanton' I have in mind a figure who is the epistemic analogue to Harry Frankfurt's

literary invention.[25] Frankfurt's wanton lacks a will toward any particular desires. The epistemic wanton lacks higher-order beliefs. She does have first-order beliefs, but no beliefs about these beliefs. She never assesses her beliefs about the world, but merely acts on whatever beliefs she finds herself possessing. She does not have, or at least does not exercise, the ability to judge that they are questionable or well-substantiated, logical or hasty, prejudiced or sincere. She essentially lacks any critical faculty. Perhaps some limited ability used to arrive at first-order beliefs might go by this name, but ruled out is any greater capacity that would give her secondary beliefs. She has no (first-personal) impetus towards revising any of the beliefs she finds herself in possession of. Notice that like Frankfurt's original wanton, our epistemic wanton seems peculiarly unautonomous, peculiarly passive. Without any second-order capabilities, the epistemic wanton seems adrift on the seas of her perceptual apparatus.

What this parallel to Frankfurt suggests is that there is a paradigm of knowledge with similarities to the reflective structure Frankfurt and others argue is the essential characteristic of freely willing human agents, indeed, of persons. Just as persons have secondary volitions which accord with or reject their primary-level desires, so they can reflectively assess and reject or endorse their beliefs. This is not to say that *all* belief undergoes conscious scrutiny. And it is a further question, even given such reflective scrutiny, what control we can exercise over our beliefs. We do not want to fall into a brute doxastic voluntarism, the suggestion that we can simply believe at will. Neither should one accept an over-intellectualized picture of belief formation. But if not all human knowledge is a type of reflective knowledge, at least some must be. The ability to

[25] Harry Frankfurt, "Freedom of the Will and the Concept of a Person," *Journal of Philosophy* 68 (January 1971): 5-20.

revise and assess, endorse or doubt, to have a variety of attitudes towards belief on occasions of reflection, seems a hallmark of *human* knowing. Beliefs which can earn the approval of reflective endorsement are of a kind that philosophers have long considered specially desirable. Why should this be important? At least in part because this seems to assure a certain kind of agency of the knowing subject. Lack of any beliefs of this certified or certifiable sort would strip a person of epistemic agency particularly important to human beings, an agency which distinguishes them from epistemic wantons.

At this point I do not want to commit to or even imply a conception of epistemic agency too closely tied to Frankfurt's account of personhood and free practical agency. Our contrast between the epistemic wanton and the fully human epistemic agent is meant to draw attention to certain characteristics of our doxastic agency—our second order attitudes and their importance to that agency, in particular. But I do not want to commit to the idea that, for example, a resounding higher-order endorsement is *constitutive* of our epistemic responsibility or agency; I want to remain agnostic, at this point, with respect to the metaphysical questions of what our agency demands of autonomy or volition. Since my purpose in looking at agency is to focus on the trouble the skeptic potentially poses for our self-conception, I do want to sketch an account that has a significant place for the first-personal point of view, and also for the idea of normative governance. So the account I'll sketch will be, in this sense, more robust than what a more purely externalist account might offer. But this just *is* to give an account that preserves and attempts to make sense of agency, rather than one that is more content to put it aside. Certainly *one* way of explaining a relatively more robust conception of epistemic agency traces a route from the human capacity for reflexivity (the fact that humans have second-order thoughts

about their first-order beliefs) to the capacity for reflective endorsement and to epistemic normativity and responsibility. Christine Korsgaard discusses how such a capability for self-reflection gives rise to associated problems of normativity:

> We human animals turn our attention on to our perceptions and desires themselves, and we are conscious *of* them. That is why we can think *about* them.
>
> And this sets us a problem no other animal has. It is the problem of the normative. For the capacity to turn our attention onto our own mental activities is also a capacity to distance ourselves from them and to call them into question. I perceive, and I find myself with a powerful impulse to believe. But I back up and bring that impulse into view and then I have a certain distance. Now the impulse doesn't dominate me and now I have a problem. Shall I believe? Is this perception really a *reason* to believe?...The reflective mind cannot settle for perception and desire, not just as such. It needs a *reason*. Otherwise, at least as long as it reflects, it cannot commit itself or go forward.[26]

Certainly it is part of our self-conception, as human epistemic agents, that we are capable of acting on reason, of acting for reasons, of believing for reasons. And the skeptic raises problems for this very basic aspect of our self-conception. In the next section I have more to say about how the skeptic does this. Here I want to expand a bit on our preliminary sketch of agency, looking some more at our higher-order beliefs and the normative practices of which they are a crucial part, pursuing a conception of agency implicit in these practices, drawing out the role of normative governance in our self-conception as epistemic agents.

For not only do we not regard ourselves as epistemic wantons, but we do not so regard most other humans. Rather, we engage with them in such ways that show we have quite the contrary assumption. Philip Pettit and Michael Smith have characterized what they call the *conversational stance* we commonly take towards others.[27] That this stance

---

[26] Christine Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996), 93. Italics are Korsgaard's.

[27] Philip Pettit and Michael Smith, "Freedom in Belief and Desire," *The Journal of Philosophy* 93 (September 1996).

describes a kind of robust agency can be seen in how they contrast their own view with Dennett's *intentional stance*[28]:

> Under the standard image, believers and desirers may be extremely rational, being well-attuned to demands of evidence, demands of consistency, and the like. They may realize almost perfectly for example, the Bayesian model of theoretical and practical coherence. But under that image, people can remain passive or mechanical subjects who harmonize and update their beliefs and desires in a more or less autonomic way...The picture of responsible believing and desiring that we associate with the conversational stance suggests a very different style of attitude formation...the subject is certainly not a mere passive or mechanical system. She does not just revise her beliefs and desires autonomically...[29]

Pettit and Smith's *conversational stance* describes the attitude we implicitly take towards another when we engage in deliberation or conversation of some intellectual manner with that other.[30] They find it significant that we do engage in such exchanges; rather than independently form all of our own beliefs, and perhaps only mechanically report on them when called to, we instead expect to be able to engage with others intellectually. This engagement reveals the expectations we have of each other and ourselves. In taking up the *conversational stance:* we assume that our conversational partner forms beliefs, and that we can both have beliefs about a common subject matter. We assume that our partner is, and sees us as, at least a potentially competent believer; we assume that if our beliefs are at odds, and if we are careful and attentive enough, we could arrive at agreement, based on good enough evidence. Pettit and Smith of course do not assume that we always expect to agree (we all know people with whom we rather expect to remain in disagreement!), but the idea is that we generally sustain confidence in the capability for explanation of who or what is right via a determination of the evidence.

---

[28] Daniel Dennett, *The Intentional Stance* (Cambridge: MIT, 1987).
[29] Pettit and Smith, "Freedom in Belief and Desire," 441-2.
[30] Or with ourselves—Pettit and Smith also suggest that we take such an attitude, such a stance towards our future or past selves, regarding such a self as a rational authority whose deliberative agency is respected, and with whom we can enter into a sort of conversation. See p. 432 on such a self-regarding conversational stance.

We expect that we can decide what we should believe based on the evidence (provided that it is adequate), and we expect that our own, and others', beliefs are responsive to such evaluations. We assume some explanation for discrepancy is available, some explanation short of the "last-resort decision...[that the other is] out of their mind and not worthy of attention: that they are not even presumptive authorities."[31] So our intellectual exchanges exhibit—and exhibit an expectation of—real engagement[32]: we regard our selves and others as presumptive authorities, as active, reasoning sources of belief and reasons for belief. Pettit and Smith also give criteria for the rationality of such a stance:

> authorizing a subject as a conversational interlocutor makes sense only if there are certain norms governing what that subject ought to believe, the subject is disposed to recognize those norms, and she is disposed to respond in the way required.[33]

In other words we assume our ability, and the ability of others, to be responsive to doxastic norms. (And that there *are* such norms and that we both assume and follow them. It wouldn't do much good, for example, to try to talk to someone about their reasons for thinking that Nader was a terrible presidential candidate if they weren't disposed to find being terrible contrary to being wonderful or getting the most votes implying not getting the least [amongst several runners].) And we assume this ability is an agency—it is not a pre-ordained rational law-conforming match, wholly passive. It is real responsiveness. We assume that we are free and responsible believers. To see another as a potential authority, we have to believe in their epistemic ability to comply with rational norms (and that there *are* rational norms):

> so far as you continue to authorize the person in conversation, you have to think that he would come around to the right belief in the event of your pressing him with the demands of

---

[31] Ibid., 430-1.
[32] Or: they *can*, and often do. Pettit and Smith of course point out that not all conversation is of this sort. Not all talking is conversation, not all conversation is intellectual exchange, and not everything that poses as intellectual exchange *is* intellectual exchange. See 430, 432.
[33] Ibid., 436.

> the evidence...you have to think that for anything that the interlocutor wrongly believes, or at least wrongly believes by available lights, he is capable of believing otherwise.[34]

Insofar as we do take up this stance toward others, we embody these assumptions. So following Pettit and Smith, then, we have a further characterization, though admittedly still quite rough and ready, of how our practices reveal assumptions people make about each other and their mutual doxastic agency. Second-order attitudes, with their normative content, show that people believe that there are things one ought and ought not to believe, and that there are better and worse, lamentable and commendable doxastic practices. But further, Pettit and Smith argue that we suppose that we are able to be responsive to these second-order attitudes, responsive to the reasons such norms embody. Skepticism forces attention to questions about such confident assumptions: are we in fact responsive to reasons in any significant, and not merely superficial, way? Are we indeed capable of normative governance? Do our doxastic practices really embody coherent and rational normative principles? We will discuss the skeptic's challenge in terms of these questions in both the next section and further, throughout the course of this dissertation study.

1.6 How the skeptic upsets our normative assumptions

Part of our self-conception is that we are capable of being responsive to reason. We are capable, though not always, but in some significant respect, of normative governance. We think this about our actions and also about our beliefs. We deliberate, we investigate, we converse and argue with others about what to believe, and on what basis. Of course this is one aspect of our cognitive economy, our cognitive practices. Not all, or even

[34] Ibid., 446.

most, of our beliefs are arrived at in such a belabored manner. But our capacity for rational and self-conscious determination is a significant part of our self-concept. Along with this conception goes a certain sort of requisite open-mindedness: to be governed or self-governing by the light of reason is to hold open the possibility that even the seemingly best founded or most entrenched of our beliefs in principle might come to require reflective scrutiny, could possibly be revoked or revised. At least this idea becomes compelling in the abstract when we focus on the idea or ideal of normative governance. If open to evidence and to reason, our beliefs are defeasible. When we think instead of certain specific beliefs, it is nearly incredible to think anything could rationally unseat them. And perhaps this really is so. Indeed, there may be sound philosophical arguments for unrevisability of certain core beliefs. I do not think that either concession substantially undoes our attraction to the general picture and the abstract belief. To accomplish this, we would need a more specific fleshing out of what goes wrong in our naïve ideas about the requirements and tenability of normative governance. Again, this merits its own discussion in chapter 3; here I hope we can at least provisionally accept this naïve assumption of rational agency—a picture of agency displayed in some basic assumption (here remaining unrefined) of our capacity for normative governance.

There are many skeptical arguments. One sort that can seem, in fairly short order, to upset the general picture of normative governance is a family of arguments building from uneliminated possibilities in combination with principles of closure and indistinguishability. Let's look at a typical Cartesian-style argument, a variation on the premises we discussed in connection with Schiffer, § 1.2 above:

(1) I don't know that I'm not a BIV.
(2) If I don't know that I'm not a BIV, then I don't know that I have hands.
So: (3) I don't know that I have hands.

Of course the BIV hypothesis is one specific skeptical hypothesis (H), and the Moorean claim "I know I have hands" is supposed to be a generic example of some fact we ordinarily think we know (O). So, more generally, the above becomes:

(1) I don't know that not-H.
(2) If I don't know that not-H, then I don't know O.
So: (3) I don't know O.

Let's recall what Schiffer said supports our acceptance of (1): "we don't count ourselves as *knowing* that the defeating etiology [BIVdom] doesn't obtain unless we have evidence for its not obtaining that goes beyond what one normally has when one takes oneself to have perceptual knowledge."[35] Schiffer seems to simply state that this is a fact about our epistemic practices, a fact about how we use and understand the epistemic concept 'knowledge'. It is hard to see whether this is so or not. *Some* people immediately accept that not being a BIV is *not* the sort of thing they know, or could know. But others do not, and may insist that they *do* know that they are not a BIV. Given that this last reaction is not merely an aberrant response, some elaboration is required if Schiffer is to be counted right about "our practices". One possibility is suggested by David Lewis. Saying that his account was meant to explain how the skeptic's argument manages to be persuasive, he briefly addresses those who persistently do *not* find it so:

> If you are still a contented fallibilist, despite my plea to hear the sceptical argument afresh, you will probably be discontented...you will insist that those far-fetched possibilities of error that we attend to at the behest of the sceptic are nevertheless possibilities we could properly have ignored. You will say that no amount of attention can, by itself, turn them into relevant alternatives.

[35] Schiffer, "Contextualist Solutions to Scepticism," 332.

> If you say this, we have reached a standoff...if you continue to find [the sceptic's argument] eminently resistible in all contexts, you have no need of any such explanation. We just disagree about the explanandum phenomenon.[36]

Lewis directs his comments to those who are not swayed by the skeptic's arguments, and so are, as Lewis says, not 'in need' of his explanation of their persuasiveness. Even when confronted with a skeptical argument, the fallibilists Lewis addresses here are *not* worried about BIVdom, they don't think they need to *ever* worry about BIVdom in order to claim to know that, e.g., they have hands. They just don't think our concept of knowledge commits them to any such *need* to worry. Whereas Lewis thinks that such a worry comes fairly easily, naturally, in epistemological contexts because when we reflect it seems as if knowledge has to be infallible.[37] Lewis calls this difference a standoff. The disagreement is, Lewis says, about the "explanandum phenomenon"—what facts need explaining.

So it is less than clear that Schiffer's claim will hold up as simply capturing the truth about our practices of knowledge ascription. Perhaps we should say, about Schiffer's claim on behalf of our practices, that *insofar as we do* say we don't know we are not BIVs, *this* is what needs accounting for or reconciling with the conflicting intuitions that lead to skeptical paradox. But then this is probably more concessive than Schiffer would like. What might further back the thought that (a tendency towards) denying knowledge in this case really does represent some significant aspect of our practices of knowledge ascription or assessment?

The thought must be that the elaboration Schiffer offers for what he calls 'our' accepting (1) draws on some significant normative principle our concept encompasses. This puts him at odds with the contented fallibilists, who presumably think this is not the

---

[36] Lewis, "Elusive Knowledge," 560-1.
[37] See ibid., 549.

case. Even on reflection, *these* epistemologists do not find themselves in skeptical paradox (at least not via the train of thought Lewis proposes, via infallibilism). They cannot find their way to a normative principle that requires them to admit (1) or (3). But Schiffer may be able to make headway, if not with these philosophers, with those less firmly decided. Even those who initially claim to know that they are not BIVs might be persuaded to change their mind on reflection, if it is the case that (1) does indeed draw on some significant normative principle implicit in our ordinary use of 'knowledge', and Schiffer can call this to attention. For if this is the case, then there is the sense that we should consistently apply this principle, even if in some instances we in fact do not. This is consonant with the skeptic's reply to the charge that he *simply* insists on impossibly high standards, insists that knowledge requires elimination of all possibilities, or complete certainty, where *normally* we do not: the skeptic will say that the fact that we do not is a *practical* fact about our everyday context, that although it is certainly true of 'our practices' that we do often ignore (due to pragmatic constraints of the everyday) many unchecked possibilities when ascribing knowledge to others, we are nonetheless committed to finding this epistemically baseless, indefensible upon reflection. The skeptic wants us to admit that knowledge precludes ignoring unchecked possibilities, that we are epistemically committed to complete checking, even if pragmatically we cannot comply. If intuitions about BIV cases really are divided, this is presumably because our intuitions differ about how to extend our epistemic concepts to this strange and remote case. What is a consistent application of our concepts, of our practices here? The skeptic can be seen as raising the question of what standards, what normative principles we do

base our ascriptions upon, of what normative principles we are reflectively committed to, given that we do believe ourselves responsive and responsible to norms.

So let's look once more at Schiffer's principle: "we don't count ourselves as *knowing* that the defeating etiology [BIVdom] doesn't obtain unless we have evidence for its not obtaining that goes beyond what one normally has when one takes oneself to have perceptual knowledge." What general principles are at work here? One candidate is an indistinguishability principle: you don't know p, if your evidence for p is the same as your evidence for q, some alternative to p. If your evidence that the butler did it is the same as your evidence that the maid did it (they were the two individuals in the victim's house at the time of his murder), then you don't know the butler did it. At the least, you'd need evidence that could favor one possibility over the other: you'd need to rule out the maid. Analogously, we can't say we know we're not BIVs because it runs afoul of the indistinguishability principle: our evidence for or against BIVdom is the same as for or against normalcy, because our experience is *ex hypothesi* the same either way.

Is this a real normative requirement? At the very least, it does seem, more generally, that it is hard to sustain the thought that one knows something once a contradictory, uneliminated possibility is introduced. As Lewis says:

> 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...I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error.' Even if you've numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong?[38]

It is hard to sustain the thought that you know, once someone points out what you haven't ruled out. All the evidence points to Jones—there are several corroborating eye-

[38] Ibid., 549-50. Italics are Lewis's.

witnesses, with good sight, in clear daylight, and no apparent motive for deception; there was the weapon seen in his possession the day before; he cannot otherwise adequately account for his whereabouts. It is an open and shut case—it is clear, you know that he did it. Then it is pointed out that Jones's hometown has the highest incidence of identical twins in the country, and you haven't found out if Jones has a twin; that the eyewitness are all white, and Jones is black, and that there is a statistically alarming rate of inadequate identification under these circumstances; that this unusual-looking knife found at the scene is actually a current best-seller amongst local gang members; that Jones is known to be a loner. Now you won't claim to know. But here you have specific reasons to withdraw your earlier claim. In any case there are always *some* uneliminated possibilities, but not always ones that seem as relevant as these. How reasonable or seemingly probable must they be, or do they only need to be bare possibilities to undermine the thought that you know? Does it make sense to admit that you have not ruled some alternative possibilities out, but still think that you nevertheless know? We know that *if* those possibilities *do* obtain, then it will be the case that you don't in fact know. Doesn't this make it sheer luck, if it turns out that you *are* right? And if it's luck that you get things right, then how is it an instance of knowledge? Isn't this paradoxical?

Here's Robert Fogelin on this issue of the effect of a remote defeater:

> Do I, for example, know my own name? This seems to me to be as sure a piece of knowledge as I possess. But perhaps, through a mix-up in the hospital, I am a changeling. I'm really Herbert Ortcutt, and the person who is called "Ortcutt" is actually RJF. These things, after all, do happen. Given this possibility, do I know my own name? I'm inclined to say that I do not. Not only that, philosophical naïfs, namely those who do not see that such an admission may lead to forlorn skepticism, tend to agree.[39]

Of course it is far from clear that we do have to rule out uneliminated possibilities in

[39] Fogelin, *Pyrrhonian Reflections*, 93.

order to know. But it can seem that we do. Especially, as Fogelin notes, before we see that this can easily lead to skepticism. I have two suggestions. The first is that the skeptic has at least succeeded in raising the issue of what normative principles our epistemic practices and our concept of knowledge do embody, and what normative principles we are committed to. We have not seen that our principles commit us to skeptical paradox, but we cannot see our way clear that they do *not*. If something like the indiscernibility principle, or a more general requirement on uneliminated possibilities, really is a normative principle inherent in our practices or concepts, one that we are reflectively inclined to accept, then we are in trouble. We never do rule out all alternative possibilities to our beliefs, not even the best of them. It is a short route to showing that such normative principles misfire and lead to skeptical paradox. But suppose it is the case that our normative principles are not really incoherent but rather we philosophically misunderstand what our principles really are. What about those 'philosophical naïfs' Fogelin mentions who are not thinking of skepticism or coherence? Insofar as 'philosophical naïfs' are inclined to accept an indiscernibility principle or an elimination requirement, then we can expect to be vulnerable to such reasoning in the course of our more ordinary epistemic lives. Then we should expect problems of incoherence or irrationality in our *ordinary* practices. At least if we are indeed capable of normative governance, of being reflectively responsive to epistemic principles or norms.

1.7 An adequate response to the normative disquiet of skepticism

If we do focus in particular on the normative difficulties suggested by skepticism, I suggest two criteria emerge for a satisfying philosophical response.

The first is an anti-insulationism: an explanation of skeptical paradox shouldn't suggest or require that we insulate our practical selves from our reflective selves. Many accounts suggest that philosophy or, more particularly, modern epistemology, creates or sustains skeptical worry. For this is to be a good response to skepticism, we need a clear sense that the philosophy or epistemology at fault is obviously distinct from merely sustained reflectivity, and removed from it. We need to be able to see that the sort of mechanisms and arguments that lead to skepticism are indeed discontinuous from reasoning we respond to in everyday life, that the philosophical misunderstandings that generate skepticism stem from theory that grows in no real way from norms or reflective ideals we respond to in real life. If this is so, then skepticism really is only an epistemologists' problem, because its subject matter is really only epistemology. There is the following problem for this type of solution: on the one hand, if the solution is too simplistic, it is not likely to take. We see how the solution is supposed to dispense with skepticism, but continue to find skepticism more compelling than the theory has resources to make sense of. But a deeper, more complicated solution of this type faces a different problem. The deeper, or more entrenched a problem skepticism is admitted to be, the less likely it seems that it is 'merely' philosophy that produced it. This kind of theory has the burden of making a convincing case that philosophy is 'artificial' and disconnected from the reflective attitudes and concerns of ordinary people.[40]

[40] Of course, there is the possibility that either the poor theory has, over time, pervasively informed ordinary ideas, or that the poor theory itself grew from poor ordinary ideas. This line of response finds philosophy rather more continuous with the everyday, but faults *both* for commitment to ideas or concepts that are optional, and serve us poorly. The possibility for this type of critical response will come for further discussion in chapter 4, where I suggest that one should meet skepticism with a *normative* argument for a non-skeptical interpretation of knowledge.

Other responses do not ascribe skepticism to optional philosophical theorizing, but do link it particularly to reflective contexts in general. Accounts that respond to skepticism by embracing a bi-perspectivalism or divide between ordinary reflectivity and practical rationality don't seem promising solutions to our problems of normative disquiet. They make it hard, if not impossible, for our reflective selves to endorse the norms of practical action we are governed by. This in turn makes trouble for the very idea of epistemic agency, and for epistemic responsibility. In the next chapter I will look in detail at epistemological contextualism, a response to skepticism of this sort. I argue that it serves to exacerbate the normative concerns raised by skepticism rather than quiet them. Chapter 3 examines, in greater detail, different facets of epistemic agency.

The second criterion for a response to skepticism that satisfies our normative concerns is that such an account make good sense of *our* epistemic concepts. This may seem like an empty requirement, or a totally obvious one. But it does seem that some accounts clearly fail to do this. And, in fact, I suggested earlier that it is less than clear how successfully this *can* be done; perhaps our concept of knowledge, for example, really is incoherent. I also suggested that it might turn out that we cannot wholly make sense of a singular concept, or that reflecting on practices we cannot sort out a concept or set of norms we are committed to. The criterion should probably be restated as the thought that we need to be able to end up with concepts we can use, that fit the roles and do the work *we*, human agents, need them to (as opposed to, for example, perfectly rational or perfectly Bayesian epistemic agents). This puts the normative question back even farther; Sally Haslanger suggests that we inquire what work our epistemic concepts do, and what values they endorse, or should:

> My concern is that because we employ our epistemic vocabulary to evaluate each other cognitively, we must undertake a normative inquiry into what is epistemically valuable, and not simply assume that "our" ordinary concept of knowledge—even when modified by recognized experts—captures what we should value. The approach I favor, then might reasonably be considered a form of immanent epistemology, but a critical or normative immanent epistemology.[41]

In chapter 4, I will look at a normative claim that the skeptic makes on behalf of our concept of knowledge, and argue that our concept need not meet this demand. Haslanger's idea that we pursue a critical epistemology is helpful, because it allows that we don't have to be definitive about what 'our' concept legislates or doesn't, but that we can refuse the skeptic by rejecting that which doesn't serve our overall epistemic values, epistemic values we endorse.

[41] Haslanger, "What Knowledge Is and Ought to Be," 467.

Chapter 2

Contextualist Solutions and Normative Perplexity

Contextualist solutions to skepticism offer a semantics for knowledge-ascribing and knowledge-denying sentences. This semantics is supposed to constitute a solution to skeptical paradox because it supposedly explains how it is that we find ourselves inclined to accept a set of statements that are mutually inconsistent. That is, we are inclined to reject the skeptic's conclusion, incredible as it is, but we are also inclined to accept her premises. Certainly, there are different versions of skeptical argument. We wonder if any of these arguments are sound, and if not, why not. But no one disputes the validity of (at least some well-constructed) skeptical argument. What is wanted is a convincing explanation of *why* we are (at least *prima facie*) persuaded by the inconsistent set of statements; the contextualist claims his semantics can do just that.

I do not think it does any such thing; I do not think the contextualist 'solution' is an adequate response to skepticism. In fact, I think the contextualist solution actually exacerbates the normative worries awakened by skepticism. One reason for this is that contextualists generally concede too much to skepticism; another is that they too little understand its real grip. In this chapter, I will discuss how contextualist solutions fail to live up to their own aims, and how they fall afoul of our requirements (on a satisfying response to skepticism) introduced in chapter 1. Our examination will show many of the problems for contextualism to be bound up with a failure to adequately consider epistemic agency; it is my intent that our study of contextualism should serve as an

example of why agency needs to be treated in conjunction with skeptical problems, a topic we will go on to consider further in chapter 3.

2.1 Skeptical arguments and contextualist targets

There are a number of treatments of skepticism that have been called 'contextualist'. Here I discuss that family of views put forward in recent years by Stewart Cohen, Keith DeRose, David Lewis.[1] For these contextualists, it is attributions (and denials) of *knowledge* that are relative to context. And while facts about the context of the *subject* (the candidate knower) will matter in assigning truth conditions to statements ascribing or denying knowledge to that subject (most obviously, whether or not the subject's belief is true), for these philosophers the context of *attribution* plays the key role in their accounts. The context of attribution is the context of the speaker (or thinker) who makes the claim about the subject (who may or may not be the same person, or in the same context). In this chapter, when I refer to contextualists and contextualism, it is these particular contextualists I intend.

Now even among this one group of contextualists, different claims are made as to what contextualism has to offer. So before looking more closely at the details of the contextualist's semantics, I want to have an initial look at the strength of claims made on behalf of contextualism (as a weapon against skepticism), and also at the type of skeptical argument addressed by each author.

---

[1] Lewis first suggested his view of these matters in a brief section of "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8 (1979): 339-359; see pp. 354-55. Lewis develops and sets out his contextualism in "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996): 549-567. Cohen's early papers are "Knowledge, Context, and Social Standards," *Synthese* 73 (1987): 3-26 and "How to be a Fallibilist," *Philosophical Perspectives* 2 (1988): 581-605. DeRose's main paper is "Solving the Skeptical Problem," *Philosophical Review* 104 (1995): 1-52. Cohen and DeRose continue discussion and defense of their views in a series of further papers, up to the present; see the bibliography (and notes below) for references to several of these more recent papers. Significantly different forms of contextualism have been propounded by Michael Williams, in *Unnatural Doubts*, and by Robert Fogelin, in *Pyrrhonian Reflections*.

2.1.1 Lewis

Like the other contextualists, David Lewis's goal is explanatory. He says that his aim is to explain how it is that skeptical arguments can be so 'irresistible', despite the fact that their conclusions are so unpalatable. The type of argument he considers is an argument from infallibilism:

> The sceptical argument is nothing new or fancy. It is just this: it 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.

Let's call the problem that Lewis is concerned with the *puzzle of infallible knowledge*. It can be represented by three claims:

Puzzle of Infallible Knowledge (PIF)

(F1) *S* knows that *P*.

(F2) If *S* knows that *P*, then it cannot be the case that there is some uneliminated possibility that not-*P*.

(F3) There exists some not-*P* possibility that *S* has not eliminated.

It is only for some *P* that the triad proves problematic. For necessary truths (F3) will be false. And of course, for any given subject *S*, there will always be many values of *P* for which (F1) is false. None of *us* knows everything. But we suppose we do know all sort of ordinary facts, and it is here that PIF gets its bite. In such cases, we confidently believe (F1) true. But we can easily be brought to acknowledge (F3); for such contingent truths, there is almost always, it seems, some uneliminated possibility of error. Lewis gives the following list of candidate defeaters, in case we get stalled: "CIA plots,

hallucinogens in the tap water, conspiracies to deceive..."[2] But our potential defeaters needn't always be *so* paranoid. I say I know where my car is at, meaning where I've parked it, although in Baltimore it is not altogether improbable that it is no longer there. But if we accept (F2), and I am unable rule out the possibility that my Civic has been stolen, then it seems I really *don't* know it's there after all. So if we find (F2) attractive, we have a version of skeptical paradox. In such cases, we are inclined to accept (F1), (F2) and (F3), but in so doing, find ourselves accepting an inconsistent trio of beliefs.

2.1.2 DeRose

Keith DeRose titled the major essay presenting his contextualism "Solving the Skeptical Problem"; he claims that "the new solution I present...*can* finally solve this perennially thorny philosophical problem...the fully articulated solution lies in the direction I point to here."[3] The specific form of skeptical argument he addresses he calls the "Argument from Ignorance." *H* is an effective skeptical hypothesis, like the bodiless brain-in-a-vat hypothesis, and *O* is some ordinary fact, ordinarily thought to be known:

The Argument From Ignorance (AI)[4]

1. I don't know that not-H.

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

C. I don't know that O.

This is the argument that we looked at briefly in chapter 1, section 6. *1* and *2* commit us to conclusion *C*; when 1 and 2 are paired with our everyday belief:

E. I know that O.

---

[2] Lewis, "Elusive Knowledge," 549.
[3] DeRose, "Solving the Skeptical Problem," 3.
[4] Ibid., 1. DeRose's formulation of this argument as a paradox is in debt to Stewart Cohen, whom I discuss just below. See Cohen, "How to be a Fallibilist," *Philosophical Perspectives* 2 (1988).

we have the argument in the form of a skeptical paradox. Although AI is similar to Lewis's PIF, AI does not mention elimination of possibilities and more explicitly draws on a closure principle for knowledge in premise 2.[5]

2.1.3 Cohen

Stewart Cohen focuses on this same argument from skeptical hypothesis[6], but has been more modest, or at least more explicit, in what he claims on behalf of his solution. He suggests that "there is no such thing as *the* problem [of skepticism], there is no such thing as *the* solution." He admits that there may be another "basis for [C]...the deductive closure argument constituted by [1 and 2 above]...is [not] the only skeptical argument...there should be nothing surprising about the fact that a response to one kind of skeptical argument does not apply to another kind of skeptical argument."[7] So while Cohen claims that his contextualism responds to AI, he does not claim that he has a blanket 'solution to skepticism'; there may be other problems left unaddressed by his theory.

Having laid out these contextualists' stated aims, we are in a position to look at their solution.

---

[5] Closure under logical entailment says that if S knows that p and p entails q, then S knows that q. Closure under known entailment, a more plausible principle, says that if S knows that p and S knows that p entails q, then S knows that q.

[6] Though Cohen always presents the argument instead as a paradox of three plausible, inconsistent claims, so instead of C, we have its negation as part of a triad:

(1) I know I have a hand,
(2) I do not know I am not a brain-in-a-vat.
(3) If I know I have a hand, then I know I am not a brain-in-a-vat.

From p. 94 of Cohen, "Contextualism Defended: Comments on Richard Feldman's 'Skeptical Problems, Contextualists Solutions,'" *Philosophical Studies* 103 (2001). See also Cohen's "How to Be a Fallibilist," 93-4; "Contextualism, Skepticism, and the Structure of Reasons," 62.

[7] Stewart Cohen, "Contextualism Defended," 96. Cohen describes the problem which Feldman raises as 'the circularity problem': "the puzzle about the structure of justification and the claim that there is no non-circular justification of our perceptual beliefs"; he admits his contextualism does not address it.

2.2 The bones of the contextualist solution

Contextualist theories are compatibilist solutions to skeptical paradox. I say this because these views allow both that what the skeptic says is true, and that what Moore, or rather a non-philosophical Moore, in a non-philosophical context, says is true. For the contextualist, it is true to say, ordinarily, that I know I have hands. But equally, when the skeptic denies that I know this, what *he* says is true. Rather than single out one of the skeptic's premises as false, the contextualist offers an argument and explanation of how we can non-paradoxically accept both the argument from ignorance *and* accept everyday knowledge claims. Using DeRose's version of the skeptical argument, this means we accept:

1) I don't know that I'm not a handless BIV (brain in a vat).
2) If I don't know that I'm not a handless BIV, then I don't know that I have hands.

C) I don't know that I have hands.

And we accept:

E) (Everyday) I do know that I have hands!

The contextualists purport to explain away our sense of paradox here; the contextualist way of reconciling the skeptic's conclusion with our workaday conviction is semantical. In short, their solution is that (C) and (E) can both be true because they do not, or do not necessarily, express contradictory propositions. They do not necessarily express contradictory propositions because knowledge claims are indexical. Contextualists hold that sentences which ascribe knowledge have as part of their content some kind of silent nod to context; part of the *meaning* of such an ascription is filled out by the context of utterance. This puts them at odds with opponents they like to call *invariantists*. Invariantists think that (given some constant assignment of content to the obviously

indexical 'I'), different utterances of knowledge ascriptions like (E) express the same proposition. Contextualists hold that the proposition expressed by (E) is incomplete and underdetermined without a specification of context of utterance. (Again, this is supposed to be the case even if we substitute some unique proper name for the indexical 'I'; the context-sensitivity we're concerned with is the alleged indexicality of what is expressed by 'knows'.) So contextualism says that (C), uttered as a conclusion to the skeptic's argument, and (E), as we might *normally* say or think it in an everyday situation (say I've just had a close call with a meat-grinder and am wondering whether shock is interfering with my perceptions), belong to different contexts. In that case they do not disagree about the truth of a single proposition, but may each express a different truth.

Of course the contextualist needs to offer motivation for his semantics aside from the utility of turning aside skeptical paradox. And he will owe an explanation of the mechanisms of shifting contexts, and some explanation of how it is that (C) and (E) *do* seem to conflict, given his semantic analysis that they really do not.

In answer to this first point, contextualists appeal to a host of examples of what we say when in non-philosophical contexts.[8] They argue that our everyday knowledge-talk displays the context-sensitivity their theory elaborates. DeRose suggests, for example, that it is an everyday phenomenon that when more importance rides on knowing some fact, this creates a more stringent context. Then the standards in place for attributing knowledge are more demanding than in other less critical, more relaxed, contexts. If a replacement liver is being flown in for a dying friend and timing is crucial, I may not be said to know the plane will be in by 6 PM, if I have just casually glanced at a schedule.

[8] DeRose, "Contextualism and Knowledge Attributions"; Cohen makes a similar point in "Contextualism and Skepticism," *Philosophical Issues* 10 (2000): 95-97.

But for a more routine rendezvous, this method might suffice. If hosts are unclear when their visitor is arriving, but want to time their dinner party appropriately, I might step in to clear things up by truly saying, "I know the flight will be here by 6 PM, there's only one flight a day from Halifax, I've checked the schedule." According to DeRose, the contexts of utterances are such that it is easier for a knowledge attribution to be true in this second case.

What context essentially determines is the *standard* which must be met for a knowledge claim to be true. This is the missing element when C) or E) is taken out of context and what is necessary in order to determine what complete proposition is expressed by an utterance of a sentence like E). The standard of evaluation is set by the context of attribution—the conversational context of the attributor. Of course there are factors other than importance of consequence serving to fix standards in a context. Cohen describes the relevant standard-setting parameters as follows: "the standards that govern a context are determined by a complicated pattern of interaction among the intentions, expectations, and presuppositions of the members of the conversational context."[9] What Cohen has in mind is Lewis's notion of conversational score. Conversational score tracks what has been said in a conversation, evolves accordingly, and dictates elements like the values of demonstratives or names, the presuppositions in play, what is assertible.[10] So the standards relevant to the truth of knowledge claims are, for contextualists, another component of conversational score. At bottom, we have a view where standards determine the truth conditions of a knowledge claim, standards are set by context, and context is determined by conversation.

---

[9] Cohen, "Contextualism Defended," 92.
[10] Lewis, "Scorekeeping in a Language Game."

2.3 A diagnosis and an error theory

The contextualist diagnosis of skepticism, then, is that the skeptic exploits the rules governing conversational score. In so doing, she creates a context where her denials of knowledge are true. The contextualist thus offers a surprisingly conciliatory diagnosis—it grants the skeptic her conclusion. But it does so only by way of robbing it of importance: skepticism, while true, is not the surprising and unsettling discovery it might seem. It is not because its truth has no tendency to impugn our ordinary knowledge; our everyday claims to knowledge stand unscathed. They are insulated from the skeptic's conclusion. This is what, instead, is the surprise discovery—studying skepticism we have not found that we do not or can not know anything in any seriously undermining sense; our ordinary knowledge claims can still stand. We have, following the contextualist, instead found that we are prone to badly misunderstand what it is that the skeptic says. Worse, we misunderstand what it is that we say.

The contextualist is, as Steven Schiffer has pointed out, committed to an error theory.[11] We ordinary folks, speaking ordinarily, in claiming to know tend not to think that what we express is tied to contextual standards. This is why we take "*S* does not know that *p*" to contradict "*S* knows that *p*"; this is why what the skeptic says seems to be directly at loggerheads with what we ordinarily think. The contextualist of course says that they can both be true in different contexts, but we never really thought this before. The contextualist is not too dismayed. Rather than take such opacity as a strike

[11] Schiffer, "Contextualist Solutions to Scepticism." Cohen points out that this is not an error theory in the sense that it regards our knowledge attributions as mostly false, or not truth-apt: "In this respect it is the antithesis of an error-theory." Cohen's right; as I commented in chapter 1 in discussing Lewis, the contextualists' methodological point of departure is that our attributions are largely true. Schiffer's nonetheless calling it an 'error theory' points to the fact that contextualists commit to the idea that we suffer a pervasive misunderstanding of our own use of 'knows that' and also commit to a corresponding potential for a type of misguided judgment when under the spell of our ability to be confused .

against his theory's plausibility, the contextualist embraces it as the explanation of why skepticism *seemed* problematic. Our lack of understanding of our own semantic usage is the explanation he needs to account for why the skeptic's argument presented us with an (apparent) paradox. According to his semantics, there is no paradox. When we take the contextualism of knowledge attributions into account, the conflict between C) and E) dissolves. But if the contextualism were readily apparent, obvious to all, we'd never have been very confused by skeptical arguments.

Schiffer questions the contextualist's semantics; he points out that the tenacious opacity necessary to his solution makes the indexicality of knowledge ascriptions unlike most other types of indexicality. With other types of indexicality, we know what's going on, we know what we mean to say. When pressed, we can draw out what's implicit in the content of what we are expressing. We seem to lack this kind of self-awareness with regard to articulating standards that the contextualist finds implicit in our knowledge claims.

> One who implicitly says that it's raining in London in uttering "it's raining' knows full well what proposition she's asserting; if articulate, she can tell you that what she meant and was implicitly stating was that it was raining in London. But no ordinary person who utters "I know that *p*', however articulate, would dream of telling you that what he meant and was implicitly stating was that he knew that *p* relative to such-and-such standard. If, for example, this ordinary guy says "I know that Placido Domingo is scheduled to sing at the Met this season' and you ask him what exactly he said, he'll tell you that what exactly he said, and meant, was that he knew that Placido Domingo was scheduled to sing at the Met this season.[12]

Talk of standards just doesn't seem relevant to what we mean to say when we claim to know some fact.

But if we don't seem to think that standards are an implicit part of the content of what we express, nevertheless sometimes we do seem aware of some sort of context-

[12] Schiffer, "Contextualist Solutions to Scepticism," 326.

variability to our use of 'know'. Playing a game of Trivial Pursuit I am considered to know the right answer if I can come up with it. I count as knowing that '1492, Columbus sailed the ocean blue'. But such a claim to knowledge, based on recollection of my 4th grade class with Mrs. Schwartzbart, might not go far in discussion with a history scholar; she will need and use different standards for her work. We know and track *this* kind of difference in contexts—if the scholar is playing Trivial Pursuit with me, surely we know and cry foul if she irritatingly tries to import her standards to the game's context. We should be able to speak more loosely.

Schiffer does admit this kind of contextual variability. He attributes it to the fact that "the penumbras of vague terms can dilate or constrict according to conversational purposes,"[13] and compares this sort of vagueness-related contextual variability of 'know' to that which we find with different uses of terms like 'flat'. And this does indeed sound just like what the contextualists had in mind all along. DeRose explicitly works from Peter Unger's early writings that assimilate 'know' to terms like 'flat'. But although Unger grouped terms like 'know', 'empty' and 'flat' together (calling them 'absolute' terms), he thought our use of these terms was a kind of hyperbole such that our flatness-ascriptions or empty-ascriptions or knowledge-ascriptions were (almost) always false.[14] Contextualists reinterpret Unger's semantics, construing his conditions of assertibility as truth conditions, preferring a semantics that counts most of our utterances true (since truth is the clearest explanation for assertibility). But Schiffer still objects that "this sort

[13] Ibid., 327.
[14] Unger, *Ignorance*.

of [vagueness-related] variability is of no use to the Contextualist...speakers are perfectly aware of when it's going on."[15]

In (an unpublished) response to this objection, DeRose disagrees, claiming that this sort of 'bamboozlement' happens all the time: "In some important sense, we don't always know what we mean, though we mean it nonetheless."[16] DeRose gives the example of a student who insists that what he means by 'know' is simply that he feels confident. Surely, this is not how the student *uses* the term, and DeRose gives the student examples in order to bring him to see this. Speakers can, in use, mark a relevant distinction without being reflectively aware of the distinction they make. Of course this is true. And like DeRose's student, one can be brought to see such facts about one's usage; this is, after all, much of the fun of reading J. L. Austin.

But this reply, so far as it has been developed, proves an unsatisfying response to Schiffer's objections. In fact, it makes one wonder if DeRose fully understands the objection Schiffer makes. For I take it that Schiffer would readily agree that speakers can in *some* cases mean something, without being able to fully articulate what it is that they mean. I don't think Schiffer would be put out with DeRose's type of example, where an individual is hard-pressed to come up with an adequate definition of a complex concept. After all, I can quite successfully mean to discuss a nuclear particle accelerator, even when at the same time I wouldn't even *recognize* one were I to accidentally come across one, so far from competent am I to give some adequate definitional explanation of just exactly what it is that I mean to be discussing. Schiffer's concern about intended meaning is *much* more specific. He takes the fact that speakers aren't aware of what the

---

[15] Schiffer, "Contextualist Solutions to Scepticism," 328.
[16] "Response to Schiffer," manuscript on-line, available from http://pantheon.yale.edu/~kd47/Response-to-Schiffer.htm. Cohen takes a similar line also—see "Contextualism Defended," 90-1.

contextualists allege is going on to pose a more particular problem for the contextualists. His objection to the contextualists' error theory comes *in concert with* his concern about what sort of account of context-sensitivity the contextualists can offer. The error theory the contextualists sign on to makes it hard to assimilate 'knows' and cognates to other already accepted and understood examples of context sensitivity. And if the contextualist is going to use a semantic analysis to address skeptical paradox, it seems he owes a plausible semantics. Here, this requires either assimilating 'knows' and cognates to already understood models of context sensitivity, or proffering a new account.

So Schiffer's objection is not: in general, to mean something by a term, one has to be able to express what one intended by use of that term. (This seems to be the straw man that DeRose addresses via his example of the student who incompetently defines what he means by 'knows'.) And it is also not: there is no contextual variability in our use of 'knows'. Schiffer has already admitted that there is some vagueness-related interest-relative variability in our usage that we tend to be well aware of, as it occurs. His objection is rather that something along the lines of awareness or intention does seem to be a requirement *in an importantly relevant class of cases*—the sorts of cases the contextualist draws on.

One standard explanation of the context-sensitivity of predicates like 'tall', 'flat', 'rich' and 'skinny' makes use of the idea of propositional content contributed by unarticulated constituents. And in cases of propositions with unarticulated constituents, it does seem that the speaker means one thing rather than another in virtue of having an intention, which could, with prodding, be specified. An example of Schiffer's we've already looked at: when Joe says "It's raining," the proposition expressed by his utterance

contains an unarticulated constituent. Joe doesn't express the proposition that raining is currently happening in some abstract or absolute sense; rather, the propositional content of his utterance contains implicit reference to some particular place. Joe, looking out the window in London, intends to express the thought that is raining in London (and expects his audience to recognize this intention, and so on[17]); London is an unarticulated constituent in the content of the proposition expressed by Joe. The unarticulated constituents used to explain the context-sensitivity of utterances like "He is rich" are intended comparison classes: neighbors, peers, Americans, and the like. So Schiffer's point is that there is an important disanalogy between the context-sensitivity at work with these predicates and with the alleged context-sensitivity of 'knows'—people aren't typically aware of intending any more complete, specifiable proposition (that articulates the standard at work) when they assert "*S* knows that *p*". It follows, we should note, that if "*S* knows that *p*" is incomplete without reference to contextual standards as the contextualists contend, it cannot be that speaker's intentions fill the gap between the incompletely specified proposition and the proposition expressed. As I suggested above, the contextualists' instead gesture towards Lewis's account of conversational score to do this work.

There is an interesting set of examples offered by Thomas Hofweber[18] on the contextualists' behalf. Hofweber argues that Schiffer overlooks an important set of cases. He groups the cases under the heading of 'hidden relativity': these are cases where,

---

[17] in the manner of Paul Grice, *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press, 1989).

[18] Thomas Hofweber, "Contextualism and the Meaning-Intention Problem," in *Cognition, Agency and Rationality*, ed. by K. Korta, E. Sosa and J. Arrozola (Dordrecht: Kluwer, 1999), 93-104. Also available on-line from http://www-personal.umich.edu/~hofweber/papers/con.pdf. Page references are to this online version.

contra Schiffer's worries, "the proposition expressed by an utterance of a speaker contains an unarticulated constituent even though the speaker is unaware that this is so and what this constituent is."[19] Hofweber's examples concern cases where empirical findings show that some property or relation is more complex that the members of some whole linguistic community have understood. He makes plausible the explanation that in such cases utterances by members of these communities are best understood as expressing propositions with unarticulated constituents of which the speakers are unaware—hidden relativity. Hofweber's first example: Prior to our discovery that motion is relative to a frame of reference, motion was thought to be absolute. Moving at *x* miles per hour was thought to be a property of an object, not a relation between an object and a frame of reference. Second example: One can assume that, at a time when a linguistic community in the Northern hemisphere had no members who had traveled to the Southern hemisphere, being winter was thought a property of certain months like January and February. Consideration of the Southern hemisphere, however, makes it apparent that being a winter month is a relation between a month and a hemisphere. Hofweber plausibly argues that the best way to make sense of these examples is to suppose that when speakers in such benighted communities utter statements like, "That cannon ball is moving at a hundred miles per hour" or "February is a winter month," hidden relativity is at work. Rather than describe the speakers as expressing incomplete or false propositions, assume hidden relativity—unarticulated constituents of the propositional content expressed, without the awareness of the speakers.

But even if we take these clever examples to tell in favor of the possibility of hidden relativity, hidden relativity of *this sort* may be of inadequate help to the contextualist.

[19] Ibid., 10-1.

Hofweber himself admits the important point that these examples have a unified content in their unarticulated constituents that makes them more benign than the hidden relativity at work under the contextualists' theory of knowledge attributions. Here, since there is a common unarticulated constituent—the Northern hemisphere, or the frame of reference of the earth—shared by an entire linguistic community, it causes no significant trouble for confused communicative intention. With the kind of context sensitivity of knowledge suggested by our contextualists, it is the case that within a language community we can misunderstand when we disagree with each other, a potentiality not readily found with Hofweber's types of examples. I'll discuss the significance of this kind of misunderstanding below, in section 2.8.

2.4 The adequacy of the diagnosis

But the important point at issue is really whether the contextualist's semantic solution is adequate to address skepticism. If there *is* an important sort of context-variability to our use of knowledge-ascriptions, does this really help? Is this putative fact diagnostically helpful, moreover, sufficient? To be truly helpful, the contextualist solution has to both convincingly explain the generation of skepticism and also convincingly explain *away* skepticism—convincingly relieve the disquiet skepticism can bring into being. So our reaction to the contextualist's bringing to light the contextuality of our knowledge-ascriptions should be somewhat like the student brought to see that he doesn't really mean 'feeling very confident' by 'know'. We should see it, and, seeing the truth of our usage, our sense of the skeptical paradox as a paradox should disappear. The AI (Argument from Ignorance—see section 2.1.2 above) should then look to us like a

fallacious argument from ambiguity, or a joke based on a cancelable implicature: a trading on a semantic brand of veiled confusion—bamboozlement.

For this to be so, we have to understand the skeptic as making an extraordinary highest-standard claim and we have to able to consider the skeptic's highest-standard denial of knowledge irrelevant to our ordinary lower-standard claims. Our agreement with the contextualist on this second point has to have a normative, not just a descriptive, dimension: we must judge that our ordinary claims do not and *should not* have to fulfill high standards to be true. Recall that the skeptic readily agrees about facts of usage—that we ordinarily do make claims to knowledge without having ruled out all alternative possibilities, that we, in this way, employ 'low standards'. The skeptic wants us to see this as *inappropriate*, that these standards themselves cannot survive scrutiny. But the contextualist says any failure to survive scrutiny is *only* due to the fact that scrutiny can change the subject. Scrutiny makes higher standards seem necessary to knowledge because it *makes* them necessary, but only by changing the very claim that is evaluated; our sense of the inappropriateness of the original claim stems from confounding contexts.

The contextualist connects what we always knew—that what seems right to say in an epistemological context is at odds with what seems right to say in an everyday context—to a more widespread pattern of interest-relative or subject-relative usage: just as what counts as flat for tennis is irrelevant to what counts as flat for skiing, so what counts as knowing for epistemology is irrelevant to what counts as knowing for train-meeting. The skeptic wants us to question whether we really know *simpliciter*. The contextualist responds, in effect, that there is no such question. I introduced considerations above (section 2.3) intended to diminish confidence that the contextualist has done enough work

on his semantic theory to convincingly assimilate 'knows' into this other class of predicates ('flat', 'rich', 'tall', and the like). I will take up this issue again in section 2.8. Aside from this semantic issue, though, there remain at least two further questions of adequacy. If a contextualist semantics of 'know' and cognates does bear out, does a contextualist semantics even show that the skeptic's conclusion is true? And if we do grant the contextualist his semantics, *and* his semantic account of the skeptic's conclusion, does this semantics offer a satisfactory response to skepticism? The first of these questions I'll take up in the following section, section 2.5. But first I want to further address the more general concerns of this latter question. Can the semantic facts adduced by the contextualists really offer up a sufficient diagnosis?

It seems unlikely that the descriptivist semantics offered by the contextualist can adequately address an important aspect of the phenomenology of skepticism: reflective discomfort.[20] Reflective discomfort stems from the press of normative imperatives that are not easily dispelled with a diagnosis of linguistic confusion, and not easily passed off as *mere* linguistic confusion. The contextualist advertised his solution as offering an explanation of a certain type: an explanation of skeptical paradox adequate for us to "proceed with confidence and with understanding to free ourselves from the trap [of skepticism]."[21] But an important part of the phenomenology of skepticism involves the alienation, from the reflective point of view, from the reasons that normally constitute our commitment to our beliefs, from seeing them as adequate, as possibly adequate. So to

---

[20] Most famously expressed by Hume (*Treatise* I iv 7): "The intense view of these manifold contradictions and imperfections in human reason has so wrought upon me, and heated my brain, that I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another....I am confounded with all these questions, and begin to fancy myself in the most deplorable condition imaginable, inviron'd with the deepest darkness, and utterly depriv'd of the use of every member and faculty...this philosophical melancholy and delirium..."

[21] DeRose, "Solving the Skeptical Problem," 3.

say that our beliefs constitute knowledge in a practical context, if not necessarily in a more reflective context, because our language, the way we use 'knowledge', makes it so, only misses the skeptic's worry, or aggravates it. To say that, on the one hand, we make a low-standard claim, but on the other, we are reflectively apt to consider a question of whether we satisfy a higher-standard claim seems to *describe* but not solve the normative problem here. What the skeptic wants to know is: why isn't that reflective context normatively dominant? If our practices of knowledge-ascription are insulationist, insulating what is truly assertible in a practical context from what is true in some more reflective context, are they rationally sustainable? The contextualist reading of the inappropriateness of claiming to know in certain reflective contexts seems to reduce the normative question that the skeptic wants to insist on to a descriptive question: what standards *do* we conform to in a given context? But there is the further question of what standards *should be* in play. The skeptic wants us to feel that we don't know because our grounds are inadequate, that they don't *amount to* knowing. This is not supposed to be merely because there is some bare possibility, some chance that we are wrong—a high standard the skeptic simply insists on. Such a simple insistence would indeed have little normative pull, little tendency to produce misgivings about our ordinary practices. A mere *demand* for high standards automatically seems as irrelevant as Stroud's example of the person who claims that there are no doctors in Manhattan, having defined doctor as someone who "has a medical degree and can cure any conceivable illness in less than two minutes".[22] The skeptic rather wants to *convince* that we don't fulfill any appropriate standard for counting as knowers. And the contextualist seems to play right into the skeptic's hands by conceding that reflectivity, itself, can *create* a high standards context

[22]Stroud, *Significance*, 40.

where we fail to know. In section 6 of chapter 1 we began to look at how the skeptic attempts to convince; in chapters 3 and 4 we will look at further elements the skeptic draws on to make her case. Until we have covered this ground, we remain agnostic about the adequacy of the skeptic's case for high standards. What is nonetheless apparent is that either this case for high standards is a good one, or it is inadequate and fails to hold up. If it fails to hold up, then the explanation of why it fails is what will be the important explanation of why skepticism (at least a skepticism based on high-standards) is not true. If it does hold up, making this case involves an explanation of why high standards are normatively appropriate, even where we do not, or can not, meet them. And even if the case for high standards holds up for some epistemological context, but not for other contexts, the explanation of why the epistemological context becomes equated with the high standard context and why the epistemological context is irrelevant to our ordinary concerns is crucial.

But the contextualist is not in the business of doing any such convincing about the appropriateness of standards, or countering it head-on, either. He rather wants us to see what the skeptic says as wholly irrelevant to ordinary knowledge-ascriptions based on the evidence he gives that our usage is context-sensitive. The contextualist wants us to see that what we really mean when we ascribe knowledge is to ascribe a context-sensitive concept: we judge that a subject's belief is reliable *enough* according to standards in play, or that her grounds (are adequate to) establish the truth of her belief,[23] according to these standards. The trouble comes with the contextualists' treatment of the skeptic's scrutiny as a mere changing of the subject to a high-standards context. In so doing, the contextualist seems to take standards as a given, the normative principles we are guided

[23] A formulation due to Robert Fogelin; see *Pyrrhonian Reflections* 19-20..

by in our practices of knowledge-ascriptions as a datum. But if an important part of skeptical worry operates by alienating us from being able to see our reasons as reasons, our selves as epistemically responsive, then contextualism does too much and too little. Too much, because treating standards as a given means the skeptic is allowed that her high standards govern (certain) reflective contexts. Too little, because by seeming to sidestep normative questions that the skeptic wants to insist on, it is then inadequate to the task of disarming the skeptic; it can even seem to exacerbate the normative concerns about rational responsiveness the skeptic raises.

2.5 Why allow the skeptic's conclusion?

The contextualist means to disarm the skeptic via his argument that what inappropriateness accompanies a claim to know in the face of unruled-out possibilities derives from operative semantic rules that are opaque to us. But it should be pointed out that this diagnosis only stands if the skeptic's conclusion expresses a true proposition. And here we find a further problem with the contextualist flattening of a normative question into a descriptive one. If it is, while doing epistemology, true to say "I do not know that I am not a BIV", this concession brings along an admission that I am unreliable with respect to the facts of BIVdom. I don't know whether I am one or not. This is supposed to be just a factor of high standards, though. So I should still think that I know according to lower standards. Just as an implicature is cancelable, so I should be able to erase bamboozlement by calling attention to context-relevant standards, and I should be able to think, in an epistemological context, that I still know according to low standards. So of course I still *believe* that I am not a BIV, and am still committed to its being *true* that I am not a BIV. But, I don't know it. This is, to use John Gibbon's

felicitous phrase, "a little Moore-paradoxical".[24] It asks that we commit to the facts about *p* at the same time as we admit that we do not know that *p*.

If one thinks one doesn't know that *p*, then one shouldn't believe that *p*, either. Richard Moran describes the implicitly evaluative, committed aspect of belief: "to have beliefs at all is for various questions to be settled...if [a person's] own belief that it is raining outside does not constitute the question being settled for her, then nothing does... to be a believer at all is to be committed to the truth of various propositions..."[25] To say that you don't know, intended merely as a concession to impossibly high infallibilist standards, is an attempt to characterize yourself wholly third-personally in a way that seems to undermine or undo the first-personal commitment that your belief embodies.[26] It is an attempt, in other words, to divorce the normative part of knowledge-ascription from its meaning, or at least to require one be able to stand apart from this normative dimension of ascriptions. It is less than clear that this concept of knowledge is one we readily recognize. Normally our knowledge-ascriptions, if they do display context-variable standards, involve an implicit commitment to those standards as appropriate. The contextualists tend to shy away from this aspect of attribution, reading standards rather more non-normatively.

Furthermore, it is this tendency, on the part of the contextualists' epistemology, that I think works exactly to exacerbate the normative concerns raised by skepticism. In chapter 1, section 5, we sketched a picture of rational agency. That sketch emphasized a conception of human knowers as appropriately subject to (implicitly evaluative) second-

---

[24] John Gibbons, "Externalism and Knowledge of the Attitudes," *The Philosophical Quarterly* 51 (2001): 15.
[25] Moran, "Self-Knowledge: Discovery, Resolution, and Undoing," *European Journal of Philosophy* 5 (1997): 147.
[26] This is an aspect of belief I discuss more thoroughly in chapter 3; see esp. section 4.

order epistemic attitudes, as agents appropriate to engage Pettit and Smith's conversational stance, as agents capable of responding to and believing for reasons. This picture contrasts human knowers with epistemic wantons. In section 6, we looked at the skeptic's upset of this conception. Where the skeptic works to alienate our reflective self from the practical self, to make our reasons feel, upon reflection, less than adequate, the contextualist willingly takes up this sort of non-normative, evaluatively-passive stance. So insofar as skepticism is a product of the taking up of such an alienated stance, contextualism can only further, not dissolve, skeptical worry. The contextualists' distancing from standards looks like a failure to inhabit and assume responsibility for the content of the epistemic judgments we make.[27] I go on in chapter 3 to discuss how I think skeptical worries are in large part worries about our capacity for epistemic responsiveness and epistemic responsibility. This tendency of the contextualist, to shuck off the normative aspect of knowledge ascription, seems to submit to the very kind of passivity that the skeptic makes us worry about. This is a point I'll return to in the course of chapter 3's discussion of agency.

Another way of getting at the issue here is to ask why the contextualist thinks that it is true to say we don't know in an epistemological or skeptical context. Why is it *appropriate* to an ascription of knowledge, *even within this context*, that the standards be so high as to be unmeetable? The contextualist will want to reply by sidestepping this question of appropriate standards. The contextualist reply will be: it is true to say we don't know in such a context, because standards *are* high, and so what is really being

[27]The idea of what it is to inhabit, rather than stand outside one's attitudes comes from Moran, who in "Impersonality, Character and Moral Expressivism," (*Journal of Philosophy* 90 [Nov. 1993]: 578-595) talks about this difference in a discussion of the limits of appropriateness of impersonality in the moral realm.

expressed by "I don't know that I have hands" is that I don't know *according to some high standard* that I have hands. What is true to say is a function of the proposition expressed, the skeptical context is a high standards context, the proposition expressed is therefore a high standard claim, and so denials of knowledge are then true. The contextualist will say: I concede that we don't know because what this means is limited by the contextual standards.

Above I suggested that normally knowledge ascriptions carry an implicit endorsement of their own standards. Saying someone knows is to endorse that their grounds establish (the truth of) their belief[28] and this is evaluative, endorsing not just the judgment that their grounds are adequate to the standards in play, but also endorsing these standards as adequate. Contextualists argue that standards are relative to context. But if this is true, isn't this because *adequacy* is relative to context? Then, standards should follow adequacy to the purpose or interests of a context; adequacy does not simply fall out of some arbitrary, given standards themselves.

To make sense of this difference, we need to consider how the contextualists define a context. What is primary for the contextualists' use of context is the idea of variable standards, and in fact it seems that this is just how the contextualists distinguish or sort contexts: one changes contexts when one changes standards. Notice the order of explanation here. It is not: different context (i.e. here we were talking about seeing someone at a baseball game, there about seeing someone at a crime scene), so different standards. Rather, it is the reverse: we count as being in a different context—where this means, e.g., a high-standards context versus a low-standards context—in virtue of the fact

[28] See Fogelin *Pyrrhonian Reflections* 19-21 for this suggestion

that the standards are changed. That is what a relevantly different context *is* for these theorists.

Epistemology is defined as a context, according to the contextualist, in terms of its (high) standards. If the kind of variability in usage contextualist semantics is predicated on is an interest or subject-relative usage (section 3), then these high standards should relate to its purpose. But the larger purpose of epistemology is to investigate and better understand our epistemic practices, knowledge and other epistemic concepts in just their everyday sense. Why would it be appropriate to securing this kind of understanding that high standards be in place? While the skeptic, and other epistemologists, have an answer, the contextualist is not interested in this question, but rather takes it as a given that the skeptic *does* succeed in driving up standards such that her conclusion is true. We will need to look more closely at the specifics of the contextualists' accounts of these mechanisms and rules governing contexts; I'll pursue this in the next section. We should here notice, though, that a context-sensitive semantics for knowledge attributions, if true, does not itself establish that the skeptic successfully or appropriately institutes a high-standards context.

In section 3, we proposed that a vagueness-account was the most likely to make sense of the contextualism we find in our practices of knowledge-ascription. Such a theory attributes a degree of contextual elasticity to our concept. But as with the vagueness of other terms, this elasticity is not limitless. If only the most anorexic of models gets counted as really skinny in the world of a fashionista, but yet a 50 pound weight-loss makes a formerly obese man, now merely doughy, get counted as "skinny!" amongst his struggling comrades, this does not make any and all nod-invoking utterances

ascribing skinniness true in *some* context. We actually have two points here about the limits of elasticity. The first is that when some concept displays interest or purpose-related contextual variation in standards, standards are variable across contexts, but within a given context, acceptable standards are limited *by those very purposes*. The second is that vague concepts may have fuzzy borders, but the variable application that might result from such vagueness doesn't imply that there is *no* limit to variability.[29]

The moral I want to draw here is that vagueness and variability are not enough to buy the contextualist the explanation he wants to offer.[30] There could be variability of standards without it being the case that there is *any* context in which the skeptic's conclusion is true. The contextualist, though, grants the skeptic her conclusion. Doing so requires a methodology that divorces the semantics of contextualism from its most promising source of support: purpose-relative or interest-relative standards. If the contextuality of standards stemmed from adequacy's relativity to purpose or interest, they would be tethered to reasons in a way that is lost when standards are instead taken as a datum read off of accepted ascription. Hooked up to whatever one might or might not

---

[29] Vagueness implies borderline cases; higher-order vagueness describes fuzziness at where the border itself lies: borderline borderline cases. Usually these go hand in hand (vague predicates exhibit higher-order vagueness). Vagueness itself should be distinguished from the subject or interest-relative variability we just flagged, although these too often go hand in hand. See Rosanna Keefe and Peter Smith, "Introduction: Theories of Vagueness," in *Vagueness: A Reader*, ed. Keefe and Smith (Cambridge: MIT Press, 1997), 1-57. I should also note that what I've said about limits to variability of a vague concept is supposed to be commonplace observation. But *how* I've put it is liable to be open to objection by some theorists of vagueness—competing theories of vagueness differ on whether vagueness is ontological or semantic, on whether e.g. vague predicates even *have* any application (see Keefe and Smith on Unger's view, p. 13). It is difficult to register the commonplace point I want to flag here, without running afoul of such objections. But my general point (that context-sensitivity of 'know' and cognates itself doesn't explain away skeptical paradox) in no way hangs in the balance. And it is the contextualist (wanting his semantics to do such work) who owes a further account of the semantics at work here.

[30] Richard Feldman offers a good example in support of this point: "The standards for 'much smarter than Einstein' might be variable—depending on how much is 'much'. But nothing meets them. The standards for 'much taller than Danny DeVito' may vary, but in virtually all realistic contexts most adults do satisfy them. Mere variability in standards does not by itself assure the sort of variability required to support contextualist replies to skepticism." See Feldman, "Skeptical Problems, Contextualist Solutions," *Philosophical Studies* 103 (2001): 67.

think of, attend to, or what someone else might or might not attend to or think of, standards seem to become too contingent, especially when the contextualist's task involves countering the normative anxieties of skepticism. The contextualist needs the skeptic's conclusion to really be true in order to have an explanation, and a solution, to skepticism. For this to be so, he has to allow a ratcheting up of standards, a methodology that stems from largely taking what is agreed upon (in a context) as true. This requires such a normative distancing, a taking of standards as given, that we land back at the problem I brought up in section 4 of this chapter. The contextualist, taking standards as given, makes them seem too contingent and rationally unsupported to sustain rational scrutiny. And this is the very worry I've alleged is at the heart of skeptical worry. (chapter 1 section 4.) To see this, and at the same time to help get a handle on some more specific commitments of the contextualists, let's consider an example.

2.6 Twelve Angry Men

After hearing the evidence, eleven jurors are convinced that they know the defendant is guilty of the murder charge for which he is being tried. They want to deliver a verdict and get home for dinner. The twelfth juror is disturbed by the complacency of the other eleven. He thinks that they should at least, together, talk it through. (Henry Fonda plays the twelfth juror in the 1957 film version of Reginald Rose's play; I'll simply call this juror 'Fonda'.) Fonda brings the group to direct greater scrutiny on the evidence and explanatory scenarios presented by the prosecutor. What had seemed airtight, under sustained scrutiny begins to seem doubtful or confusing. The group, which had initially been so convinced of the defendant's guilt that they didn't even feel the need to discuss it, ends up voting to deliver a not-guilty verdict. It is clearly not the case, however, that

they are at this point convinced of the defendant's innocence. Rather, they are no longer sure enough of his guilt.

One way to describe what has happened fits nicely with the contextualist's account. What has happened is that the jurors have, through sustained reflection, brought into play possibilities that they had previously not considered. They made relevant possibilities that they had previously ignored—that the witnesses were lying, or that their testimony was inconsistent, that the defendant's possession of a knife was explicable without the assumption that it was the murder weapon. At first, the jurors thought they knew the defendant was guilty. Each (save Fonda) thought he himself knew, and also thought the others knew. They were willing to vote to convict, to deliver a guilty verdict. Reflection on defeaters—not-ruled-out alternatives to the beliefs that constituted their evidence—served to raise the standards for knowledge. Then, though they perhaps still believed that the defendant was guilty, or that it was more probable than not, they no longer thought that they knew he was.

It might be objected that this is not a good explanation of this case. It is not a good explanation because it is not true that the standard for knowledge actually changed during the jury's deliberations. The standard was, and remained the same: reasonable doubt. The explanation for the eleven members changing their votes is not a change in standards, but rather the introduction of further evidence. When Fonda had the jurors act out, for example, the old man moving from his bed to the door, or estimate the time a subway train would take to pass the building, their so doing produced new information, new evidence. It was this new evidence that changed the jurors' minds, not any change

in the standard they applied in evaluating whether or not they knew the defendant was guilty.

If you object thus, you are likely not a contextualist. Because for the contextualist 'reasonable doubt' has to be as contextual as 'knowledge' is. What counts as knowledge is inextricably intertwined with what counts as reasonable to doubt; if what counts as knowledge is relative to context, then what is counted as reasonable to doubt is relative to context. Has the context changed in this case? Could the jurors really change contexts just by the direction of their conversation? Since the relevant notion of changing contexts here is one of the standards in play, the standards relevant to knowledge ascription, the question is whether reflection or deliberation alone can operatively effect such a change. And the contextualist is committed to this possibility. The contextualist will not want to say that the jurors changed their minds based on new evidence and not on a change in standards, because deliberation *always* produces 'new evidence' in the form of new beliefs (even in the limiting case where deliberation only serves to confirm the first order belief, it has produced 'new' evidence in the form of a confirming second order belief). Where the new evidence is in the form of a belief that formerly ignored possibilities are relevant, this seems to be just what the contextualist claims *is* a change in standards. So the explanatory line that says the jurors did not change standards because they *would* have, even early on, admitted the relevance of the new evidence had they then considered it fails to convince the contextualist. According to her, this fact just gives a description of how the deliberative process in this case can function to change standards. The jurors have been persuaded by Fonda to feel the gaps in the prosecutor's case, to feel that they are relevant, and that they therefore cannot count themselves as meeting the standard for

knowing the defendant is guilty. Initially, eleven were content to say they knew; at the movie's end, all twelve say that they do not. Initially they considered the prosecutor's case sufficient to prove guilt, later they do not.

Juror Six lends additional plausibility to the contextualist's account that what has transpired is a change in standards, in one of the film's better-known lines:

> Well, I'm not used to supposin'. I'm just a workin' man. My boss does all the supposin' -- but I'll try one. Supposin' you talk us all out of this and, uh, the kid really did knife his father?[31]

Juror Six's comments underline the fact that it really was Fonda, and Fonda's perseverance and conviction that influenced the other jurors and the eventual outcome. He 'talked them out of it', out of the guilty verdict. He convinced them that they didn't know. But of course the viewer is not meant to see Fonda as a bully. He is more the voice of justice, chipping at and exposing the prejudices of the jury's angry men. Although the viewer may share Juror Six's disquiet, this is a disquiet appropriate to the gravity of the jury's task of deciding the fate of another person, disquiet particularly appropriate in light of the human fallibilities the play calls attention to. The viewer is meant to feel Fonda has pushed the jury in the right direction. If the jurors raised their standards, this was appropriate because the young defendant's very life was at stake. Eleven of the twelve jurors had been frighteningly complacent.

But this contextualist interpretation of the film has trouble adequately accounting for this basic point: the jury's final standard is appropriate, its initial laxity is not. Consider the same jury without Fonda. Presumably such a jury would have summarily issued a guilty verdict. Their standards for ascribing knowledge—to themselves and to one another—are lower. It is natural to think their judgment, and their standards, are

[31] From the adapted screenplay by Reginald Rose, "Twelve Angry Men," 1957.

mistaken. Given what they know, they do not know the defendant is guilty. The contextualist says that knowing has a contextual component—so that to judge whether a knowledge-ascribing sentence is true or not, we need to complete the proposition by filling in the standards of the ascriber's context. Instead of knowing simpliciter, we have high-standards knowing, low-standards knowing and infinitely many standards in between. In the case at hand, we have Fonda-knowing and Angry Jury-knowing (I'll call this 'angry-knowing' for short), with different standards for each. Of course I am (and the viewer in general is) in a context where I, as an assessor, am attending to the possibilities that Fonda raised, and consider them relevant, so I subscribe to Fonda standards and not angry standards. So I am interested in whether the jury Fonda-knows, not whether the jury angry-knows. But I can also imagine that no one had ever attended to the Fonda possibilities and that in the angry jury they maintained angry standards and they angry-knew. They never Fonda-knew, but this was not what they asked themselves when they considered whether they really knew if the defendant was a murderer. They were wondering if they angry-knew, and they did. I want to say that even in this case angry-knowing isn't relevant to knowing,. They may think it is, but it really isn't. Knowing is Fonda-knowing.

But to insist on this point, for the contextualist, I have to be mistaking my context with the angry context, missing the point that the question of whether the jury Fonda-knows is distinct from the question of whether the jury angry-knows. I have the mistaken meta-belief that contradictory sounding utterances about angry-knowledge and Fonda-knowledge are in conflict although they needn't be. But this reply misplaces the concern. The concern is that the jury should be concerned with Fonda-knowledge, not with angry-

knowledge. The Fonda standards are appropriate to their charge, the angry standards are not. If whether or not they will *say* they know the defendant is guilty depends on the standard in play and if the Fonda-standard is only in play on juries lucky enough to contain Fonda, then whether or not they say they know is far too contingent to satisfy.

The opponent of contextualism will also say: there are not many different propositions that can be expressed by "I, juror *n*, know that the defendant is guilty." There is *one* relevant proposition because there is one relevant standard. The most natural reading of what has transpired with the jury is not that they first endorsed the proposition that they angry-knew, and they later quite consistently endorsed the proposition that they failed to Fonda-know the defendant's guilt. Rather, the angry eleven *disagree* with their former selves. There was a common propositional content at issue all along. At first they would have endorsed the proposition expressed by "I know that the defendant is guilty" and later they deny the truth of *that same proposition*.

Do the contextualists' views have the resources to answer these objections? We have two distinct lines of inquiry here. First, if contextualists build standards into the content of knowledge claims, do they have a plausible account of disagreements that seem to be disagreements about appropriate standards themselves? And second, if they allow the skeptic a pass on her high standards (within context), what prevents standards from correspondingly bottoming out? That is, if a high-anxiety or highly imaginative context that considers skeptical possibilities makes knowledge claims into high-standards claims, what prevents an analogously complacent context from making knowledge claims so low standard that knowledge effectively reduces to mere accepted-as-assertible true

belief? I want to first pursue this last question. To answer it, we'll have a closer look at Lewis's contextualism.

2.7 Rescuing standards: Lewis's account

Lewis is aware of the problem of 'cheap knowledge'. This is why his definition of knowledge includes reference to what is *properly* ignored, not simply what *is* ignored:

> *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.[32]

He has three prohibitive rules, outlining what may not properly be ignored[33]:

> PR1) The *Rule of Actuality* says that what is actually the case may not be ignored.
>
> PR2) The *Rule of Belief* says that what possibility the subject believes to be actual may not be ignored. Additionally, "neither is [a possibility properly ignored if it is] one that he [the subject] ought to believe to obtain—one that evidence and arguments justify him in believing—whether or not he does so believe."[34]
>
> PR3) The *Rule of Resemblance* says that if some possibility 'saliently resembles' another that may not be ignored due to PR1) or PR2), then it may not be ignored, either.

These three prohibitive rules prevent cheap knowledge. They prevent angry-knowledge in the angry-jury from properly counting as knowledge. The additional clause to the Rule of Belief prohibits this. Even if it is not the case that any of Fonda's alternative explanatory hypotheses were themselves convincing enough that they ought to have been believed, a further addendum to Lewis's Rule of Belief covers the case: "A possibility may not be properly ignored if the subject gives it, or ought to give it, a degree of belief that is sufficiently high...How high is 'sufficiently high'? That may depend on how

---

[32] Lewis, "Elusive Knowledge," 554.
[33] See Ibid., 554-556.
[34] Ibid., 555.

much is at stake."[35]

So why did I ever suggest that the contextualist accounts are so permissive, distancing themselves from normativity by reading standards from whatever is accepted as a standard in a context? Lewis's prohibitive rules (those we've so far considered) clearly seem to belie this picture. I made the suggestion because a crucial aspect of Lewis's anti-skeptical contextualism really *does* stem from such a permissive distancing. The prohibitive rules do not show this distancing from normativity but we still need to look at Lewis's permissive rules, and one final prohibitive rule. With regard to such normative questions, there is in an interesting asymmetry in Lewis's account.[36] His *Rule of Conservativism* says to

> suppose that those around us normally do ignore certain possibilities, and it is common knowledge that they do...then—again, very defeasibly!—these generally ignored possibilities may properly be ignored. We are permitted, defeasibly, to adopt the usual and mutually expected presuppositions of those around us.

So maybe the angry jurors, *sans* Fonda, would be able to ignore the mutually ignored possibilities after all? No, presumably the defeasibility clause is built in so that, in such cases, the Rule of Belief trumps the Rule of Conservativism and the jurors are not permitted to ignore Fonda's possibilities, even without his mentioning of them. But once Fonda is in there, the *Rule of Attention* also applies:

> When we say that a possibility *is* properly ignored, we mean exactly that; we do not mean that it *could have been* properly ignored. Accordingly, a possibility not ignored at all is *ipso facto* not properly ignored. What is and what is not being ignored is a feature of the particular conversational context. No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative.[37]

---

[35] Ibid.
[36] Michael Williams notes this asymmetry amongst Lewis's rules, remarking: "Although the point of introducing presupposition-rules was to prevent knowledge-by-*fiat*, the Rule of Attention provides for ignorance at will." See "Contextualism, Externalism and Epistemic Standards," 15.
[37] Lewis, "Elusive Knowledge," 559.

The Rule of Attention is a further prohibitive rule. Lewis says it is a triviality, but it is hardly that. It is what gives him a solution to skepticism, by way of an explanation of how standards rise in certain reflective contexts. Given the Rule of Attention, Fonda-possibilities are not properly ignored once brought into play, and the relevant standard for knowledge-ascription clearly becomes, via this rule, the Fonda-standard. And this, of course, is precisely how Lewis accounts for the skeptical conclusion's truth within its context. Skeptical possibilities, once attended to, may no longer be properly ignored:

> Do some epistemology. Let your fantasies rip. Find uneliminated possibilities of error everywhere. Now that you are attending to them, just as I told you to, you are no longer ignoring them, properly or otherwise. So you have landed in a context with an enormously rich domain of potential counter-examples to ascriptions of knowledge. In such an extraordinary context, with such a rich domain, it never can happen (well, hardly ever) that an ascription of knowledge is true. Not an ascription of knowledge to yourself (either to your present self or to your earlier self, untainted by epistemology); and not an ascription of knowledge to others. That is how epistemology destroys knowledge. But is does so only temporarily. The pastime of epistemology does not plunge us forevermore into its special context. We can still do a lot of proper ignoring, a lot of knowing, and a lot of true ascribing of knowledge to ourselves and others, the rest of the time.[38]

Attention is really a thin concept for Lewis, which makes this rule a very strong one. It really does seem like a mere mention of a possibility serves to make it relevant (at least temporarily), and improper to ignore. He says, of the sort of epistemology that investigates what may or may not be properly ignored: "to investigate the ignoring of them was *ipso facto* not to ignore them...that is how knowledge is elusive. Examine it, and straightway it vanishes." Lewis also says, "if you bring some hitherto ignored possibility to our attention, then straightway we are not ignoring it at all, so *a fortiori* we are not properly ignoring it. How can this alteration of our conversational state be undone? If you are persistent, perhaps it cannot be undone..."[39] Lewis goes on to say that an effort at ignoring can, in time, turn into real ignoring, if all parties are cooperative.

---

[38] Ibid.
[39] Ibid., 560.

There is some escape clause to the Rule of Attention. But little enough to make knowledge seem far too elusive. We are vulnerable not just to skepticism but to paranoia and to epistemic bullying. It is not only contemplating evil demons, or brains-in-vats, but thinking of most *any* specific uneliminated possibility, however actually remote, in a specific case, that can serve to rob one of knowledge. And this is just what Lewis says:

> Even if S himself is neither sceptical nor an epistemologist, he may yet be clever at thinking up far-fetched possibilities that are uneliminated by his evidence. . . even if S's idle cleverness does not lead S himself to draw sceptical conclusions, it nevertheless limits the knowledge that we can truly ascribe to him when attentive to his state of mind. More simple: his cleverness limits his knowledge. He would have known more, had he been less imaginative.[40]

But if this is so, the worrier and the well-informed cannot know so much, as long as they are worrying or thinking of all the uneliminated possibilities. If the possibility of home-invasion is ever close to mind, I can never know that my loved ones are safe at home, no matter how remote the possibility, no matter how obsessive or paranoid my thinking. If I think it a bare possibility that some neighbor lying in wait could just reach his hand far enough into the garbage chute to intercept, as it passes, the bag I have just thrown down there, then I cannot know it has reached the basement, even if it has.[41]

The contemplative are hit particularly hard. Reflective endorsement, except for the particularly uncreative, seems an impossible ideal. We cannot know so much, so long as we are reflective. At least this is certainly the case with a contextualism that allows a mechanism similar to Lewis's Rule of Attention. This rule makes it hard to understand reflectivity as having any normative epistemic role. Reflection seems as liable to lead us

---

[40] Ibid., 561.
[41] This is a variation on an example of Sosa's, "Skepticism and Contextualism," *Philosophical Issues* 10 (2000): 13-4.

astray, to make relevant crazy possibilities, and rob us of what would otherwise be knowledge, as to lead to a better set of beliefs.

If skepticism is in large part, as I have suggested, bound up with normative worries about our ability for reflective responsiveness, then Lewis's elusive knowledge only seems to aggravate, not effectively explain away these concerns. Instead it seems to institute a bifurcation between our reflective and practical selves, for on his account it is precisely too much reflection and too much attention that destroys knowledge. And Lewis even admits this bifurcation by way of an example of what he calls 'compartmentalization':

> A compartmentalized thinker who indulges in epistemology can destroy his knowledge, yet retain it as well. Imagine two epistemologists on a bushwalk. As they walk, they talk. They mention all manner of far-fetched possibilities of error. By attending to these normally ignored possibilities they destroy the knowledge they normally possess. Yet all the while they know where they are...the compartment in charge of philosophical talk attends to far-fetched possibilities of error. The compartment in charge of navigation does not. One compartment loses its knowledge, the other retains its knowledge.[42]

And if Lewis's attempt to show skepticism true relies on an elastic Rule of Attention that makes knowledge seem far too contingent an achievement, his other prohibitive rules are correspondingly strict. Strict enough to avoid cheap knowledge by design, they sneak normativity back into the contextualist's story. But they end up too strict to be explanatorily useful—the Rule of Resemblance and the clauses of the Rule of Belief and Actuality together prohibit almost any proper ignoring. So how do we get ordinary knowledge in ordinary contexts? By what Lewis admits is an ad hoc maneuver:

> any other possibility *W* that is likewise uneliminated by the subject's evidence thereby resembles actuality in one salient respect: namely, in respect of the subject's evidence. ...plainly, we dare not apply the rules of Actuality and Resemblance to conclude that any such *W* is a relevant alternative—that would be capitulation to scepticism . the Rule of Resemblance was never meant to apply to *this* resemblance![43]

---

[42] Lewis, "Elusive Knowledge," 565.
[43] Ibid., 556.

Lewis leaves normativity out of his Rule of Attention to allow it to ratchet up standards to skeptical levels, when proper ignoring reduces to virtually nil. But he cannot leave normative clauses out of his rules altogether to avoid the corresponding problem of cheap knowledge. He builds the right results into his rules to get ascriptions of knowledge we normally consider true to come out so under his rules, but then has to escape skeptical consequences in ordinary contexts by an ad hoc exception. This is not an explanatory account that can possibly respond to the type of skeptical worry we introduced in chapter 1.

It does avoid the problem of 'cheap knowledge'. How Lewis does this has revealed an interesting asymmetry in his (and other contextualist) account(s). For these theorists, contextual standards are integral to mapping truth conditions onto knowledge attributions, integral to determining when it is true to say "S knows that p". Lewis's Rule of Attention purports to record the fact that (mere) mention of not-*p* possibilities creates standards such that said possibilities must be ruled out for S to truly count as knowing p. The status of this Rule is wholly descriptive—that is, this rule for proper ignoring I take it is supposed to be like rules describing conversational scorekeeping: Lewis is describing how knowledge attribution works, by enumerating the several rules we adhere to in our practices of knowledge ascription. (We might think of the Rule of Attention as a standards-tracking component of conversational score.) When we look at the addendum to the Rule of Belief that prevents 'cheap knowledge', we see a normative component enter the rule: "A possibility may not be properly ignored if the subject gives it, *or ought to* give it, a degree of belief that is sufficiently high..."[44] Why the normative clause

[44] Ibid., 555. My italics.

preventing cheap knowledge, but no normative clause preventing inappropriately rarified skeptical standards?

The answer has to be: this is the way our concept works. The asymmetry is in our concept of knowledge: we let standards get limitlessly high, we don't let them get too low. Recall my discussion in chapter 1 of Lewis's Moorean strategy of taking (most of) our attributions to be true, sorting out our concept of knowledge from the rules that can be summarized in service of making consistent sense of those attributions. Mere mention of remote defeaters and we are liable to reject knowledge claims—standards have been raised, we fail to know. Little scrutiny and consideration of defeaters and we are liable to accept them—standards are low, we do know. But this last quick conclusion—standards are low, we do know—has an important caveat. Standards can't get too low—if we ought to give some attention to some possibility, we can't ignore it—we *don't* really know, even if, conversationally, all are agreeably complacent. Knowledge can't be that cheap.

2.8 What we really know, what we really *want* to know

Presumably knowledge can't be that cheap because we have to retain the idea that knowledge is an achievement, some kind of epistemic good over and above true belief. If standards could bottom out, we'd lose that aspect of our concept. It is certainly right for any account of knowledge to preserve this basic attribute. But Lewis arguably gives up on other significant aspects of our concept. Recall that Lewis treats knowledge as a matter of having eliminated possibilities—all possibilities (save those properly ignored; these properly ignored possibilities are what prove contextual). He explains:

> Why have a notion of knowledge that works in the way I described?...it is one of the messy short-cuts—like satisficing, like having indeterminate degrees of belief—that we resort to because we are not smart enough to live up to really high, perfectly Bayesian, standards of rationality. You cannot maintain a record of exactly which possibilities you have eliminated so far, much as you might like to...ascriptions of knowledge to yourself or others are a very sloppy way of conveying very incomplete information about the elimination of possibilities.[45]

I think that this is actually a fairly surprising view of knowledge ascriptions. The picture is one where knowledge ascriptions are wholly descriptive—saying *S* knows is 'sloppy shorthand' in lieu of more precise information about the possibilities that *S* has eliminated. Such a picture leaves out any essential function for an evaluative component of knowledge-ascription. But it is also troubling in concert with the contextualist idea that our use of 'knowledge' has a persistent opaqueness to it. Remember that Lewis's picture is of a concept that allows unspoken proper ignoring, but no eyes-wide-open, conscious ignoring. The other contextualists all share the idea that we are prone to overlook the context-sensitivity of our concept; that is, after all, why we take the skeptic's true conclusion to be at odds with the truth of everyday knowledge claims (section 3 above). But if contextualism is true, and to the extent we rather *think* we apply unitary standards or otherwise misunderstand what property it is that we ascribe when we attribute knowledge, this practice is potentially quite misplaced. Our use of the shorthand concept potentially misfires. We would be better off if we could instead pass on the perfect list of eliminated possibilities Lewis imagines. Then we would have both more accurate information, *and* a clearer view of the information we are acting upon. The perfect practice would be more transparent than the sloppy shorthand.

But what if our ascriptions of knowledge are *not* merely incompetent shorthand in lieu of more *specific* information, but rather an equally important *assessment* (that, e.g.,

---

[45] Ibid., 563.

the knower's grounds are sufficient to establish the truth of his claim)[46]? Lewis neglects any evaluative component of knowledge-ascription, the importance of ascription as judgment. If I have a witness at trial, even after I have asked all the relevant details, and gleaned all the specific information I can from her, I will still want to know what she considers herself to *know*, given these facts. We take up an attitude towards ourselves as knowers with respect to certain facts; we take up this attitude towards others. This attitude has a robustness and importance beyond that we would have towards a concept we truly treated as mere sloppy shorthand. It is an important task of this dissertation project as a whole to further characterize this robustness and importance. One place we look is at the role of ascribing (or refusing to ascribe) knowledge to oneself. In stressing the role of second-order (and higher-order) attitudes in comprising our epistemic agency, I stress the role of knowledge-ascriptions as prescriptive judgments. More generally, by virtue of the positive epistemic status conferred upon knowers, ascribing knowledge confers license—to action (based on such knowledge), to belief. Of course, if, as Lewis suggests, knowledge as a concept is fairly superficial, there is not much at stake with the threat of skepticism.[47] But I don't think this *can* be right because of the very important role knowledge plays for us by way of license. If contextualism is right, there is a troubling lack of transparency to this theory's account of our practices. *Our* concept of knowledge, as it is entrenched in our practices, is not one we often *see* as a short-cut

[46] This particular formulation is suggested (though it is only part of his analysis of knowledge; he also has an epistemic responsibility clause) by Fogelin, *Pyrrhonian Reflections*. My point here, though, is intended to be more general, and not wedded to Fogelin's specific analysis of the content of an ascriber's assessment.

[47] Or at least a knowledge-based skepticism; see chapter 4, section 1 for discussion of knowledge-based skepticism.

replacement for Bayesian standards.[48] I do think that Lewis's disregard for the concept of knowledge suggests why he rests easier with his contextualist analysis than the rest of us should.[49] At the *normative* level, he is not so far from Schiffer's conclusion that we are stuck with an incoherent concept.[50] Lewis's idea that our concept *is* infallibilist[51] drives his contextualism and necessitates the opacity that threatens normative incoherence. Of course, his semantics allows him to say that we *do* in fact have a consistent use of knowledge sentences, despite initial appearances to the contrary suggested by conflicting standards we apply at different times, suggested by the trap the skeptic is able to lead us into. We have consistency because we have, as it were, a plurality of concepts, each indexed to context. Or if not a plurality of concepts, some other specification of indexicality which explains contextuality. Though this saves consistency in the truth conditions which govern our knowledge attributions, it does this at the cost of a usage which is readily apparent to its users.

---

[48] Perhaps except when it comes to very specific domains of study, as in science. Lewis says "if you doubt that the word 'know' bears any real load in science or in metaphysics, I partly agree. The serious business of science has to do not with knowledge *per se*; but rather, with the elimination of possibilities...and with changes that one's belief system would (or might or should) undergo under the impact of such eliminations." He is perhaps right about this. But all of life is not science, and our concept of knowledge serves purposes mostly other than those of science. I will argue that one such purpose is linked to our interest in intentional action and intentional, reason-giving, explanations. It would not be surprising that a concept most useful for "the serious business of science" would be different than what is most useful for more everyday purposes (equally serious business, but indeed of a different sort). Lewis makes this case when he admits that science doesn't really use the concept of knowledge. But this, too, is debatable—one could argue that theory construction relies on a more evaluative and committal, less neutral standpoint than Lewis's eliminated-possibilities-tracking.

[49] A comment of Bernard Williams with regard to our concept of responsibility could well apply to the choices we have confronting a concept of knowledge that Lewis finds to be a messy shortcut: "...reflection [on our concept of responsibility] can go only in one of two directions: either in the direction of saying that responsible agency is a fairly superficial concept, which has a limited use in harmonizing what happens, or else that it is not a superficial concept, but that it cannot ultimately be purified." Williams, "Moral Luck," *Proceedings of the Aristotelian Society* S.V. 50 (1976). I am suggesting that Lewis chooses this first path with regard to our concept of knowledge, but that the second might be more adequate (in view of the importance knowledge has for us).

[50] See Schiffer, "Contextualist Solutions to Scepticism," 330.

[51] See section 2.1.1 above.

This opacity suggest several related problems. Each of these problems is significant and deserves more extended exploration than I will accomplish here. One problem is suggested by considerations we first considered in section 3, when we looked at potential difficulties for the contextualist's semantics. There, we discussed Schiffer's objections to coupling the contextualist's semantics with his error theory. Schiffer's argument was that in order for 'know' to be context-sensitive in a manner appropriate to solve skeptical paradox, it had to be problematically different from context-sensitive terms like 'tall' or 'flat'. It had to exhibit what Hofweber calls 'hidden relativity'. This being the case, it is obvious that speakers' intentions cannot fill the role of articulating the unspoken constituents of propositional content. The contextualists instead use features of conversational context to fix the missing component. And in order to generate a semantics that secures the skeptic's conclusion, it must be the case that these standard-fixing features don't just flow from the normative notion of purpose or interest-related adequacy—otherwise, we'd need a further explanation of how and why epistemological contexts normatively require high standards, a question the contextualists do not pursue. The contextualists relinquish the idea that there is a prior notion of appropriateness of standards aside from the relatively thin (and relatively malleable) notion of appropriateness as dictated by (conversational) context.[52] And this seems particularly troubling if reflection alone can indeed change the context to make it more demanding. Then we have the idea that standards are different for the reflective than the nonreflective, or those who manage to more often sustain lenient intellectual contexts

[52]Although we should keep in mind the asymmetries we looked at in discussion of Lewis just above. The contextualists do have a notion of appropriateness when it comes to how vulnerable our beliefs can be, at minimum, to count as knowledge; the malleability of standards comes in the open-ended possibility of *raising* standards ever higher.

than those who do not. There is no endorsement of, for example, the edicts of a particularly reflective context as 'correct' beyond their specific context, in normative governance over standards and judgment. In fact, for Lewis, such contexts often prove irrelevant to what we really and more ordinarily want to say about knowledge and who has it. This puts his view in decided tension with a picture of autonomous rational agency that give a specially important role to particularly reflective contexts. (I further discuss this picture of agency in chapter 3.) It also makes it hard to see the concept of knowledge as a fully useful concept. Raised standards, or introduced possibilities seem to force us to be engaged and concerned with issues (the truth of high-standards claims) that are not necessarily relevant to the interests or purposes at hand in contemplating and assessing our or another's knowledge. And worse, we lose track of the fact of this irrelevance. So, wanting to judge whether I can fairly and reasonable be said to know of the accused's guilt, joined with a persistently imaginative, persuasive and vocal interlocutor, according to the contextualist I can only truthfully conclude that I do not know (for standards will remain persistently high). If we concede to the contextualist his semantics, we concede this is the right answer. But in such a case it is then certainly the wrong question. It is not what we wanted to know.

We pointed out a fundamental asymmetry in Lewis's account, shared by the other contextualists. They allow that the skeptic can create a high-standards context where the truth of knowledge claims requires meeting that high standard. But they don't allow any cheap-knowledge contexts. The contextualists find this asymmetry in our concept of knowledge. Lewis is explicit that this is because our concept is an infallibilist one. Yet there seems an exact parallel between the irrelevance of angry-knowing and high-

standards *not*-knowing. They don't really bear on our interests. So neither is what we mean, or *want to* mean, when we ask ourselves if we know. That question is a fundamentally normative question, related to the license conferred by ascribing knowledge.

We first noticed the tendency of contextualist views to overlook the evaluative aspect of knowledge ascription in section 5 of this chapter. There we noticed that it is unclear what a contextualist has to say about how we should think of self-knowledge, knowledge of one's own attitudes. It is supposed to be the case, for the contextualists, that doing epistemology, and considering far-flung defeaters as I am now doing, it becomes true for me to say "I do not know that I'm not a brain in a vat." But then reading the contextualist account of our concept of knowledge, I see that this is really just a factor of the high standards in play; so I should still think that some ordinary-standards claims are not undermined in their own contexts. I should still think that I know according to ordinary standards, but don't know according to the currently elevated standards. But as we noted in section 5, to say that you don't know, intended as a concession to standards you don't even necessarily normatively endorse is a strange sort of attempt to characterize yourself wholly third-personally without occupying any first-personal, commitive standpoint. For presumably what is particular to first-personal knowledge claims is that they are avowals, not merely reports. That is, they are primarily expressive, rather than attributive, the outflow of what Richard Moran has called a 'deliberative stance'; they express a view of the *world*, and a commitment to that view, and are not just a descriptive report on one's *self* (that one knows, or doesn't). Part of the content of that view, expressed by the statement "I don't know that I have hands," might

be that one is unreliable with respect to the facts about one's hands or that one's belief in one's handedness is not sufficiently grounded. Put in *these* ways, though, such thoughts are not particularly first-personal; someone else could describe me in precisely these ways. The statement "I don't know that I have hands" *also* expresses my thought that I am explicitly *not* committed to the belief that I have hands, or that I am not fully entitled to the belief insofar as I have it. In a high standards context that I recognize as such (and recognize not to be in conflict with my more everyday beliefs): am I committed or not committed to these beliefs? Should I be? These questions, fundamental to whether I will conceive of myself as having knowledge, don't seem to be separable from whether it is 'true to say' (because of standards) that I have it or not; commitment is a fundamental aspect of knowledge (and belief).

Moran puts a fundamental point about my commitments to my beliefs this way: "as I conceive of myself as a rational agent, my awareness of my belief is awareness of my commitment to its truth, a commitment to something that transcends any description of my psychological state."[53] To think that I know that p is first to think about the world; it requires that I am able to inhabit my belief in a way that is not to self-ascribe, report or otherwise third-personally attribute that attitude or mental state. The contextualist uses the idea of shifting standards to cleave the skeptic's conclusion from its significance. But insofar as I conceive of my beliefs as rational, and conceive of myself as a rational deliberator, I cannot just separate my endorsement of my beliefs, their normative and prescriptive content, from their descriptive content. So I can't coherently see myself as just knowing according to low standards, not knowing according to high standards. That

[53] Moran, *Authority and Estrangement: An Essay on Self-Knowledge* (Princeton: Princeton University Press, 2001), 84. The line of thought developed here indebted to Moran—see n. 27 above, and further discussion of these issues in chapter 3 below.

would be simultaneously to license, and not, to inhabit, and not, commitment to my beliefs. These first-personal issues I will take up again, in chapter 3's discussion of epistemic agency, but they are also particularly complex. In covering this territory my object is to give a sense of the character and significance of this problem, but more primarily, it is the limited goal of showing that there *is* a problem here. For I think it is a serious one for the contextualist view; moreover, it is a serious problem for any similar epistemological theory, any theory that requires, as I said in first chapter, an insulationist or non-integrative view of knowledge.

We are now in a position to see that the questions left open in section 3 and section 6 converge. In section 3, we left the issue of how the alleged hidden relativity of knowledge attributions is different, and more problematic than the other cases of hidden relativity that Hofweber describes. In section 6, we left the question of whether contextualists have a plausible account of disagreements that seem to be disagreements about appropriate standards themselves, given that they build standards into the very content of knowledge claims. Both issues stem from the contextualists' commitment to opacity—the idea that we fail to fully understand our concept of knowledge insofar as we misunderstand how we use 'knowledge'. The questions 'how do we understand ourselves and each other?', and 'how should we understand when we agree or disagree?' are both related to how we disambiguate content, and the problems that stem from the contextualists doing so in a way that relies on attributing an error theory to speakers' understandings of their own utterances.

In the examples of hidden relativity that Hofweber describes, there is a common unarticulated constituent shared within an entire linguistic community. So ignorance that

there *is* some unarticulated constituent, that there *is* some hidden relativity at work, does not make for confused communicative intention. If I'm talking with my neighbor over a seed catalog about the coming summer months, neither of us is liable to be confused about whether the other makes a claim about summer months within this hemisphere, a claim about the Southern hemisphere, or some purportedly absolutist claim about the non-relativity of seasons. If Hofweber is right about hidden relativity, it is correct to say (in his cases where, *ex hypothesi*, I'm not aware of relativity) that I am not aware of the full content of what I express with the utterance that "July is a summer month." Insofar as I don't recognize the relativity of my utterance, I don't recognize the full content of what I claim. But even if we accept hidden relativity as the most plausible semantic story, and admit this consequence, the contextualist has further work to do if he is to assimilate *his* semantic story about knowledge attributions to these cases. Part of what makes hidden relativity a plausible semantic story in Hofweber's cases is *because* (even while it does postulate content of which the speaker is unaware) it doesn't make trouble for communicative intention; this semantics doesn't make trouble for our intuitive notions about what the speaker was using the sentence to communicate. With the contextualist construal of knowledge claims, things are otherwise.

In Hofweber's cases, we can easily rely on facts about the real nature of motion or seasons to give content to what is claimed by utterances about motion or seasons. This seems as plausible as Kripke's explanations of how we can successfully refer via causal communicative chains, without some accurate and unique mediating descriptive representation. We can, in some cases (as we alluded to in discussing DeRose's reply to Schiffer) successfully talk about something without understanding fully what it is we talk

about. Each component of content needn't derive from some specific mediating communicative intention. In Hofweber's cases, the work of completing content can be done by context and not by some specific communicative intention (to refer to the Northern rather than Southern hemisphere, or any hemisphere at all; to refer to the frame of reference of the Earth) since such completion makes sense of our *overall* communicative intentions without problem. Even though I don't think of the relativity of seasons to hemisphere and so don't have the specific communicative intention to convey the thought that 'July is a summer month *in the Northern Hemisphere*,' I do intend my utterance to convey some true and relevant thought about *seasons*, relevant to my neighborly conversation about planting our gardens (which, are, as a matter of fact, here in the Northern hemisphere).

Not only is it the case that we can have content without a (fully) corresponding communicative intention, we can also have content in the face of *contrary* communicative intention. Pure indexicals are cashed out by functions from context to content,[54] and so take their values independent of communicative intention, allowing for such a mismatch. For example, Rip Van Winkle might wake up to say "I fell asleep yesterday," intending to express the thought that he fell asleep on July 3, 1766.[55] But since Rip has been asleep for 20 years, his utterance expresses the proposition that he fell asleep on July 2, 1786. 'Yesterday' picks out the day before the occasion of utterance; given the meaning of 'yesterday' and the context, its designation is automatic, and does

[54] See David Kaplan, "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals," in *Themes from Kaplan*, ed. Joseph Almog, John Perry and Howard Wettstein (New York: Oxford University Press, 1989), 481-614.
[55] This example is John Perry's; see his "Indexicals and Demonstratives," in *A Companion to the Philosophy of Language*, ed. Bob Hale and Crispin Wright (Oxford: Blackwell Publishers, 1997), 595.

not depend on the intentions of the speaker (or any associated demonstration) to pick out some given day.

Other indexicals *do* depend on something more; the designation of demonstratives like 'that car' or 'that man' aren't fully automatic but depend on demonstration or salience or directing intentions for completion. The speaker's intention is relevant here to the content of the proposition expressed by his utterance; there is no completing function from merely objective features of the context to semantic values .

Now our contextualists allow that "the standards [for knowledge] that govern a context are determined by a complicated pattern of interaction among the intentions, expectations, and presuppositions of the members of the conversational context."[56] So it turns out to be a tight line the contextualists have to walk here. While their solution to skepticism requires that (insofar as we are taken in by skeptical paradox) we are mistaken about the content of our utterances of knowledge ascription, the content of these utterances does itself depend on features of our communicative intentions. The content of our knowledge-ascribing utterances is filled out (as a matter of fact) by operative semantic rules that depend on our thought and talk, our communicative intentions, our conversational context, but we fail to see this. Our failure to see this means that we can systematically misunderstand what it is that we are talking about, and which knowledge-ascriptions are inconsistent. And unlike Hofweber's cases, here we do have potential problems for significantly confused communicative intention. I want to assess whether I can fairly and reasonable be said to know of the accused's guilt, but because of my persistently imaginative flights of fancy, I can only truthfully conclude that I do not know *according to the persistently high standards my imaginings have brought into play.* I

[56] Cohen, "Contextualism Defended," 92.

confuse the content of my conclusion as addressing the question I started out with. I say there is, in these cases, "significantly confused communicative intention": there is the serious possibility not just that speakers don't grasp the *full* content of their utterances, but moreover that what the speakers intuitively use the sentence they utter to communicate comes apart from the semantic interpretation it is given by the theory.

This is an important difference. Part of what makes Hofweber's postulation of 'hidden relativity' in the cases he discusses so much as plausible is the way it saves the phenomena. Rather than attribute some incomplete proposition (without truth value) or false claims to all the members of some society prior to knowledge of hemisphere-relative seasons, we preserve the truth of their claims *and* make sense of their overall communicative intentions by way of hidden relativity. And, as we said, part of what makes hidden relativity plausible here is that the speakers were talking about *seasons*. They lacked some relevant non-semantic knowledge about seasons; since the 'missing' content is consistent across the linguistic community, there is no problem for misascribed content (between speakers, or on the part of a self-reflective thinker or speaker) that leads to serious communicative misunderstanding. The contextualists' reliance on 'hidden relativity' is otherwise. It sets up the real possibility for confused communicative intention in a more serious way.

In the case of DeRose's student, and in Hofweber's cases, it is intuitive to the *speakers* that the semantic story we propose is true to their communicative intentions insofar as they can be convinced that the semantic story spells out what they meant all along. Even though DeRose's student at first claims that when he says "I know that p," he intends to claim just that he is very confident that p, the student can be convinced (by

further consideration, further cases, further reflection) that this is *not* really what he means. He *never* really meant that. When Hofweber's benighted subjects learn about the relativity of motion, they will recognize that insofar as they meant to talk about *motion* they did implicitly discuss a framework-relative concept; even if they didn't know that all along, they meant *that* (motion) all along. Here there is a substantial sense in which the semantic story remains true to the overall communicative intentions of the speakers. But can we expect the same for the contextualist's semantic story about knowledge attributions? Not, I think, when it comes to high-standards contexts. In the kinds of cases where reflection alone, or mere consideration of remote possibilities drives up the standards, the content that is attributed to speakers by the contextualists' semantic story threatens to come apart from speakers' own abiding intuitions of what they were using their words to communicate. In our case of deliberation about the accused's guilt, I want to know whether I know in an action-licensing sense, not merely in some arbitrarily rarified standards sense. The semantic story the contextualists' give in these cases is at odds with our own sense of our communicative intentions when what we supposedly say we know, or fail to know, is not what we wanted to know.

Contextualism sets up the real possibility for confused communicative intention in a serious way, serious enough to undermine its plausibility as a semantics. It allows too much scope for serious communicative misunderstanding, and self-misunderstanding, to a degree that begins to threaten our rationality. If this is so, it hardly seems like an *alternative* to skepticism, let alone a solution to it. At the very least, these issues point up the need for a more in-depth discussion of these topics on the part of the contextualists.

There is one final point to be brought into the discussion of the troubles the opacity of the contextualists' semantics creates for its plausibility. This is the problem of disagreements that seem to be disagreements about appropriate standards themselves. The contextualist solution to AI assimilates the seeming inconsistency of "I know I have hands" and "I don't know I'm not a BIV and I don't know I have hands," to the seeming inconsistency of Tweedle-dee's "I'm hungry" and Tweedle-dum's "I'm not hungry." We clear up this silly argument by filling in the indexical content to show that nobody's really in disagreement about some single proposition here. But the skeptic wants to argue about standards themselves. And it seems likely that arguments about whether some person knows that *p* or not, are, in general, often arguments about standards. So specifying standards won't clear up the argument; it won't settle the issue of which standards are normative, relevant, appropriate. In *some* cases we might clear up a seeming disagreement between us by further spelling out standards of application that govern our use of a term: you say "Spud Webb is short," I say "Spud Webb is tall!" If you mean he's a short basketball player, and I mean he's still a tall guy, we don't really disagree. But if what we're really arguing about is his suitability for NBA play (suppose we're deciding whether we want to keep him on our team), what we really argue about is how short is too short for our team. Knowledge is like this—we wonder about adequacy of grounds; wondering about adequacy is a question of appropriateness (of standards), not confusion about which standards we invoke. Skepticism is like this—a standard-specifying knowledge-claim won't clear up the argument. Being completely clued-in to the semantic content of my knowledge claims won't make skepticism go away.[57]

[57] So the contextualist needs to give further attention to just *how* the contextualist semantics, if true, is supposed to clear up disagreements. In cases of hidden relativity that Hofweber discusses, this is fairly

The contextualists, in general, offer too tight a circle of explanation to serve the purpose they want to address. Standards are high, so skepticism is true. But skepticism is only true because standards are high. The contextualists really never give an adequate account of what needs the most attention in the first place: what, if anything, makes high standards at all relevant to ordinary knowledge? What *reason* is there to think they are demanded?

Lewis stresses the thought that there seems to be an inappropriateness to claiming to know in the face of unruled-out possibilities. With the other contextualists, he hopes his

---

clear. Being semantically clued-in clears up disagreements that are like the "Spudd Web is short/tall" or "I [Tweedle-dee/Tweedle-dum] am/am not hungry" cases. For example, suppose I've lived in a (Northern hemisphere) community where the relativity of seasons to hemisphere is so far undiscovered. Still ignorant of the hemisphere-relativity of seasons, I suddenly find myself in Australia, conversing with some Aussies. They insist that July is a winter month. I disagree. What are we to make of this exchange? We might, at first say that that I mean that *July is a summer month in the Northern hemisphere* and they mean that *July is a summer month in the Southern hemisphere*. But if I look around, see a lot of snow, experience a piercing wind and realize I've woefully underpacked, and say, "This can't be July! It's winter!", clearly I am now talking about the Southern hemisphere, where I find myself. But I am *also* thinking about my past experience in the Northern hemisphere. So what are we *now* to say about the content of my claim, "July is a summer month"? Does it express the thought that 1) July is a summer month in the Northern hemisphere 2) July is a summer month in the Southern hemisphere 3) July is a summer month in some absolute sense? or instead 4) an incomplete proposition? Of course, 1) is the option that continues to attribute some true belief to me. If I express 2) or 3) I express some false proposition. But whatever we want to say about the content of my utterances in this new circumstance, what is clear is that my confusion (about my disagreement with the Aussies) will be cleared up if someone *explains* to me the nature of seasonality. Then I will see which claims are genuinely in conflict, and which are not, and I won't get caught up in Tweedle-dee, Tweedle-dum arguments. The contextualist will want to say that his explanation of the context-relative standards that govern the truth conditions of knowledge attributions will do the same for us (particularly with respect to skepticism). But with regard to knowledge claims it is hard to see explicit acceptance of context-relative standards to clear up disagreements in a similar way. This is so only if this type of skeptical problem arises *only* from the semantic confusion the contextualist alleges is going on. Our attachment to an action-licensing concept of knowledge seems evident when we consider the types of disagreements that *can't* be cleared up by some standard-specifying knowledge-claim; standards themselves are often at issue. Insofar as the contextualist treats AI like a Tweedle-dee, Tweedle-dum argument, and insofar as the contextualist does seem to take seriously the idea that high-standards *do* govern certain reflective contexts, then the contextualist is committed to the idea that different standards are action-licensing in reflective and ordinary contexts. Or at least if it is the case that knowledge is centrally action-licensing as I argue; I say more about this idea in chapters 3 and 4. But this conclusion, that different standards are action-licensing in reflective and ordinary contexts, just restates the skeptical problem. So although I think the real problem here is the incompatibility of the contextualists' concept of knowledge and the normative character of knowledge, at the very least the contextualists need to say more about *how* the contextualist semantics is supposed to fully dissolve AI-style arguments for an agent who accepts contextualism.

semantic analysis is enough to diagnose and undermine the bite of this route to skeptical conviction. But I don't see that evidence about usage will be enough to undo all the skeptic has to say in service of such convincing or that it will be enough to fully disarm his ability to produce the conviction that we really do not know in just the sense we (meant to) express in everyday usage. I think there is more by way of deep-rooted internalist intuitions about knowledge that are not addressed by marshalling semantic evidence. I will have more to say about such internalist intuitions and their sources in chapters 3 and 4; one important source I argue for is the skeptic's (and our own) misunderstanding of the ability to rationally self-govern. But even without making this case, doubts can be raised about the sufficiency of the contextualist diagnosis.

Chapter 3

Responsibility, Agency, Authority and Authorship

In the last chapter, we looked at contextualist diagnoses of skepticism. I argued that this type of contextualism problematically divorces itself from important normative issues at the heart of skepticism. I argued that the contextualist theories fail as a solution to skepticism, both on their own terms and in terms of the desiderata we set forth in chapter 1. In this chapter, as we look more closely at epistemic agency, we will be in a better position to see how the normative dissatisfactions with contextualism we've already examined relate to contextualism's conception of agency. For one purpose in looking at contextualism was to use it as an example of a generally inadequate response to skepticism, an inadequacy related to a failure to take epistemic agency seriously enough. In this chapter, I will elaborate on the type of agency at stake, in order to better understand its importance.

3.1 Human knowledge and reflective agency

Some accounts of knowledge, like some accounts of practical agency, stress the specifically human, while others are more content to acknowledge a strong continuity between this epistemic property we are concerned to attribute to each other and the capacities of other animals, or even the type of discriminatory ability evidenced by nonliving things like light sensors, thermostats or smart valves. Obviously, theorists of this latter stripe are not much concerned to offer a very robust role or characterization of epistemic agency, and in this sense offer a thinner conception of what human knowledge amounts to.

I have already suggested that skepticism makes trouble for knowledge by way of making trouble for agency and a view of our selves as rationally responsive. One way out is to offer a response to skepticism that rather ignores these agency-based problems, instead offering up a skeptic-proof concept of knowledge that sets these particular worries about agency aside. I will discuss such an option in chapter 4. I argue that maintaining a sense of ourselves as rational agents is importantly what it is to be non-skeptical; an answer to skepticism that does not wholly assuage or adequately address these issues of agency will fail to be an adequate response to skepticism. Not a failure in the task of answering the skeptic in terms she'll accept (what James Pryor calls the "ambitious anti-skeptical project"[1]), and not a failure in the methodological task of constructing an epistemology that has the resources to handle skeptical objections (described in chapter 1 by MMV theorists as skepticism's real import, a methodological spur to adequate theorizing), but in the task of answering the skeptic in ourselves.

Ernest Sosa draws a distinction between two species of knowledge which he calls 'animal knowledge' and 'reflective knowledge'. Humans can have both, but as Sosa's terminology suggests, a capacity for the latter type is the more distinctive achievement:

> One has *animal knowledge* about one's environment, one's past, and one's own experience if one's judgements and beliefs about these are direct responses to their impact - e.g., through perception or memory - with little or no benefit of reflection or understanding.
>
> One has *reflective knowledge* if one's judgement or belief manifests not only such direct response to the fact known but also understanding of its place in a wider whole that includes one's belief and knowledge of it and how these come about.[2]

Sosa elsewhere contrasts human knowledge with what he calls 'servo-mechanic' knowledge:

---

[1] James Pryor, "The Skeptic and the Dogmatist," *Nous* 34 (2000): 517.
[2] Ernest Sosa, "Knowledge and Intellectual Virtue," in *Knowledge In Perspective: Selected Essays in Epistemology* (Cambridge: Cambridge University Press, 1991), 240.

> ...there is a sense in which even a supermarket door 'knows' when someone approaches, and in which a heating system 'knows' when the temperature in a room rises above a certain setting...human knowledge is on a higher plane of sophistication, however, precisely because of its enhanced coherence and comprehensiveness and its capacity to satisfy self-reflective curiosity.[3]

Sosa points to this "capacity to satisfy self-reflective curiosity" as a distinguishing mark of a specifically human knowledge. But, as I began to argue in chapter 1, the reflexive potential of human thought enables more than an ability to understand the "wider whole" of how one came to believe some particular, or to achieve a broad and coherent set of beliefs. Reflexivity also gives structure to a form of agency particular to fully-functioning adult human beings. In chapter 1, we considered a figure we called the 'epistemic wanton'. The epistemic wanton does not have second-order attitudes towards her first-order beliefs, and never exercises or displays the ability to reflectively assess her beliefs. Humans do have such second-order attitudes. In lacking these second-order (and any higher-order) attitudes, the epistemic wanton lacks an important kind of rational agency, for she will not ever be motivated by such assessments or by reflection to further investigate or cease to investigate, to abandon belief or confidently endorse it. We human epistemic agents seem capable of a kind of intentional activity, therefore, that the wanton lacks. We have reasons to act and motivations that the wanton does not; we can act on the basis of these reasons.

One way of explaining a relatively more robust conception of epistemic agency traces a route from this human capacity for reflexivity (the fact that humans have second-order thoughts about their first-order beliefs) to the capacity for reflective endorsement and to problems of epistemic normativity and responsibility. Human epistemic agents can have epistemic responsibility, since we have the capacity to evaluate our own beliefs

---

[3] Sosa, "Epistemology Naturalized," in *Knowledge in Perspective*, 95.

and to be motivated by those assessments—motivated to further investigation, further deliberation, to suspension or deferral, to endorsement and confident action. This picture of epistemic responsibility is closely tied to being the kind of agents capable of engaging in the *conversational stance* we introduced in chapter 1. Seeing each other as participants in the manner necessary to support the conversational stance means that we don't just see each other (or ourselves) as passive with respect to our beliefs, but potentially active, reasoning sources of belief. Our beliefs are responsive to doxastic norms since we are responsive to doxastic norms. We control beliefs to the extent that we see them as responsive to our rational assessments and deliberations, and since we can be motivated to pursue further information, further avenues of inquiry, further lines of reasoning, which we can choose to pursue or not. In these respects we can see ourselves as active (or potentially active) with respect to our belief, and not merely saddled by whatever beliefs we happen to have, subject to beliefs we find ourselves with. It is in this sense that we are free and responsible with respect to our beliefs.

This picture already finds itself opponents, for there is a stalwart tradition in epistemology of arguing that we cannot be responsible for our beliefs, since we do not have the kind of control over our beliefs necessary to support responsibility. The guiding thought behind this view is that belief is not voluntary. We can not will beliefs, and this ability to believe at will is necessary to support the kind of responsibility we want to attribute to one another, and perhaps necessary, even, to support the meaning of deontological epistemic judgments and language.

3.2 Arguments against control

Many philosophers have claimed that we cannot properly be held responsible for our beliefs. Drawing on the dictum 'ought implies can', their basic idea is that we lack the requisite freedom to believe: we lack the kind of freedom that would buttress responsibility.[4]

Believing at will

It doesn't seem to be the case that doxastic voluntarism is true—we don't have the ability to simply believe at will. I cannot believe anything I want, just because I might want to. Even if I think it might help me pay better attention, I cannot simply will myself to believe that the dreary speaker before me is the world's brightest philosopher. I can't will myself to believe that four and five is ten or eleven, that the earth is held up by a turtle, or that Andrew Lloyd Webber really is a witty composer.

Belief aims to represent the world. It has a 'direction of fit' that is mind to world. This idea of direction of fit contrasts belief with desire; belief's job of mapping the world is opposed to desire's task of shaping the world. John Searle describes the difference as a difference in "responsibility for fitting":

> Because it is the responsibility of the belief to match an independently existing state of affairs in the world, we can say that the belief has the *mind-to-world direction of fit.* It is the task of the belief, as part of the mind, to represent or fit an independently existing reality...[5]

This can explain why I can't just will myself to believe anything at all. Such an inability is not some kind of metaphysical lack or psychological block on my part, but is conceptually part of what it is to believe, an aspect of the character of the attitude of

---

[4]See Alvin Plantinga, *Warrant: The Current Debate* (New York: Oxford University Press, 1993), Chapter One; William Alston "The Deontological Conception of Epistemic Justification," *Philosophical Perspectives* 2 (1988): 257-99.

[5] John Searle, *Rationality in Action* (Cambridge: MIT Press, 2001), 37.

belief. Since I believe about a world independent of me (although I can also have beliefs about myself, but we'll leave aside such complications for the time being), and since to believe is to believe to be true, I cannot just *decide* to believe against how I know things to be. I cannot just decide to believe *anything at all*, because this wouldn't be compatible with the mind-to-world direction of fit. Insofar as belief is such that it intrinsically aims to fit, I could not decide to *believe* regardless of the way I find the world to be. Believing is conceptually too closely tied to finding or judging the world to be a certain way to leave room for such independence. It is an attitude of regarding a proposition as true.[6] I could hope or wish or imagine the world differently, but not simply believe differently in the face of the facts I genuinely take as facts.[7] When Hume noted this fact (that we cannot just decide to believe anything at will) in the *Treatise,*[8] he makes this inability sound like a purely contingent psychological fact, as if this inability were comparable to not being able to stop digestion or blood flow whenever one wanted to.[9] But I have

---

[6] This formulation is David Velleman's, from his "Introduction", *The Possibility of Practical Reasoning* (Oxford: Oxford University Press, 2000), 17. But as John Searle points out (in *Rationality in Action* [Cambridge: MIT Press, 2001], 36), one must take care not to regard a *proposition* as the object of the attitude--of course the worldly state of affairs is the object of my belief, the proposition gives the content of my belief.

[7] This is not to say that I could not have some motive other than truth come to be at odds with the motive to believe the truth, and that I might desire to preferentially serve that other motive. Recent studies have shown that those who are rather more optimistic and less realistic about the opinion others have of them are happier. If I want to be happy (more than I want to know the truth) I have a motive not to believe what is true (in terms of others' accurate opinion of me). So I could pursue a project of self-deception of this sort, seeking to acquire false beliefs. But the fact that this is a project and that it is self-deception means that I cannot *simply* be unmoved by facts that I recognize. There is also the matter that people's beliefs can be amazingly impervious in the face of facts, as if they were indifferent to them: how else could one think that slaves aren't people, that women are inherently inferior...the psychological phenomena at work are as fascinating as they are numerous and complex, and I can't begin to do them justice within the scope of this chapter, or this dissertation project. It is in light of such complications that I said one cannot *simply* believe differently in the face of facts *taken as such*. In the end, I do think these and other possibilities for irrationality are quite important to mapping out our concept of belief, how reason can motivate belief or fail to, and the role of our agency in all this. Here it suffices that there is a basic connection between belief and taking true that is a useful starting point.

[8] See Hume, *A Treatise of Human Nature*, I iii 7-8.

[9] Bernard Williams points this out in "Deciding to Believe" *Problems of the Self* (Cambridge: Cambridge University Press, 1973), 148.

already suggested that this inability is, more saliently, *conceptual*, an aspect of the attitude that belief is, part of the grammar of belief.

Bernard Williams made this link to the conceptual content of 'belief' in his well-known 1970 article "Deciding to Believe." Williams's brief, but influential argument:

> If I could acquire a belief at will, I could acquire it whether it was true or not; moreover I would know that I could acquire it whether it was true or not. If in full consciousness I could will to acquire a 'belief' irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e. as something purporting to represent reality...I could not then, in full consciousness, regard this as a belief of mine, i.e. something I take to be true, and also know that I acquired it at will. With regard to no belief could I know...that I had acquired it at will. But if I can acquire beliefs at will, I must know that I am able to do this...[10]

I want to look more closely at Williams's argument. I can agree with some, but not all, of what he says; Williams commits to a stronger thesis than the conceptual link I have wanted to bring out.

I have endorsed the idea in this passage's second sentence, the belief-aims-at-truth thesis:

BAT) One cannot consciously will a belief irrespective of its truth.

BAT rules out willing to believe in the flatness of the earth or that 2+2=5. BAT itself can be broken down into two ideas: a basic conceptual thesis about belief and a relational thesis about the will and belief. (I want to break apart BAT in this way because it will be useful to have this material at hand when we go on to look at Williams's general view of epistemic agency.) The basic conceptual thesis about belief is the (rough) idea we discussed above, aided by Searle's notion of direction of fit: since belief is "something

[10] Bernard Williams, "Deciding to Believe," *Problems of the Self* (Cambridge: Cambridge University Press, 1973), 148.

purporting to represent reality,"[11] as an attitude it cannot float freely from the way the world is. At most basic the conceptual claim is that to believe is to believe true.[12]

CT_{belief}) Beliefs (qua beliefs) aim at truth.

The conceptual thesis CT_{belief} becomes the basis for BAT. BAT itself is a thesis about the will's relation to belief. So BAT can also seem to be a truth about the scope of the will. We can see this second idea if we bring it out explicitly:

CT_{will}) The will is such that one cannot consciously will a belief irrespective of its truth.

This sounds slightly strange. It sounds strange because it seems to imply the Humean thought that this is some kind of contingent, psychological inability to control our attitudes, whereas the relational thesis of BAT really stems from the conceptual thesis about belief articulated by CT_{belief}. BAT, although it gives a limitation on our ability to will, really stems from a conceptual thesis about belief, a thesis about the type of attitude belief is, not a conceptual thesis about the will. CT_{will} might be okay as a conceptual *conclusion* about the will drawn from CT_{belief} and BAT, but it can sound confusing taken as a simple conceptual truth about the will. Williams does want to draw a conclusion about the scope of the will and about our agency with respect to belief, about the extent to which belief acquisition is up to us. In fact, he endorses a stronger claim about our lack of agency than CT_{will}. So he might well welcome the confusion potentially engendered

[11] Of course this very example brings out the necessity of more careful handling: 2+2=4 is not well described as a belief "purporting to represent reality" unless one is a mathematical Platonist. The basic conceptual truth Williams is reaching for and that I want to endorse should not be beholden to strong metaphysical views. BAT itself may be less problematic than this gloss on what supports it. But I go on below to take up how BAT can easily be conflated with more theory-laden theses.

[12] Williams says that 'to believe *p* is to believe *p* is true" is part of what he means by the slogan 'beliefs aim at truth'. See his discussion pp. 136-7.

by CT_{will}—inference to some *intrinsic* limitation of will—because he does embrace a strong thesis about the will's limits, about the voluntariness of belief.

Williams endorses the stronger claim, also on conceptual grounds, that one cannot acquire *any* belief by means of the will, even a proposition one believes *true*. He endorses:

Doxastic involuntarism$_1$: One cannot acquire beliefs by will.

Williams arrives at this involuntarism by way of the first claim in the passage I've quoted, the conditional: "if I could acquire a belief at will, I could acquire it whether it was true or not." Beyond the thought about belief's conceptual connection to truth, there is a general thought here about the voluntariness of belief. Williams's claim about 'believing at will' is presumably that having such an ability would be a matter of having basic voluntary control[13]: it would be something we could intend to do and then just do, just like that, like snapping your fingers. I don't think anyone thinks believing is like *that*. (Insofar as believing *is* something that seems to *just happen*, it doesn't seem to be something *I* do, or will; it is not something I have, in these cases, intended.) Williams's point about our incapacity to will belief is also meant to be a more general endorsement of doxastic involuntarism. He says his general topic is: "how far, if at all, believing something can be related to decision and will"[14]; and later that:

> ...there is not much room for deciding to believe. We might well think that beliefs were things which we, as it were, found we had...although we could decide whether to express these or not. In general one feels this must be on the right track.
>
> It fits in with the picture offered by Hume of belief as a passive phenomenon, something that happens to us.[15]

---

[13] Alston, "The Deontological Conception of Epistemic Justification," 263.
[14] Williams, 136.
[15] Williams, 147-8.

So Williams also seems to endorse:

Doxastic involuntarism$_2$: Belief acquisition is not up to us.

I see no reason to accept *this* broader view; in fact, we still have reason to question even doxastic involuntarism$_1$.

Doxastic involuntarism$_1$ relies on CT_{belief}, but must also rely on further assumptions. Embodied in this claim is a very particular view of the will, a view which I think we have little obligation to accept. It is a view of will I call the *lawless will*. The idea of the lawless will is that willing is a wholly unconstrained activity, that willing as an act is inherently unfettered. And rationality itself counts as constraint. This is why Williams believes "if I could acquire a belief at will, I could acquire it whether it was true or not." And Williams in turn believes doxastic involuntarism$_1$ because he thinks that *willing* belief conceptually requires willing irrespective of truth. He moves from the BAT thesis via his assumption of the lawless will to doxastic involuntarism$_1$.

Descartes had a completely different view of the will in the *Meditations*. For Descartes, willing is not incompatible with the dictates of rationality. On the contrary, proper willing for Descartes follows the dictates of rationality. Error arises when the will judges beyond that which it properly understands. The best use of the will, that which avoids error, is to keep it within bounds of true understanding. So for Descartes, the idea of the will being bound by understanding is not a real limitation on the will, but a recommended form of exercised restraint:

> ...since the will is much wider in its range and compass than the understanding, I do not restrain it within the same bounds, but extend it also to things which I do not understand: and as the will is of itself indifferent to these, it easily falls into error and sin, and chooses the evil for the good, or the false for the true.

> ...as often as I so restrain my will within the limits of my knowledge that it forms no judgment except on matters which are clearly and distinctly represented to it by the understanding, I can never be deceived...Nor have I only learned to-day what I should avoid in order that I may not err, but also how I should act in order to arrive at a knowledge of the truth; for without doubt I shall arrive at this end if I devote my attention sufficiently to those things which I perfectly understand; and if I separate from these that which I only understand confusedly and with obscurity. (Meditation IV)

Above I quoted Williams endorsing the idea that belief is a passive phenomenon. There is no *willing* belief because there is no room for us to choose anything. The idea of the lawless will I have attributed to Williams seems to follow from the thought that if rationality (or anything else) dictates choice, there is no conceptual room for will to enter the picture. It is superfluous. While I don't think that Williams provides any argument for his assumptions about the will, I also don't think that the notion of the will is particularly *helpful* for the questions about control and agency that Williams, more generally, addresses—the issues of doxastic involuntarism$_2$. It is really this thesis of doxastic involuntarism$_2$, Williams's more *general* picture that I am concerned to combat. And I don't think that it is necessary, if we are to defend some capacity for or variety of control over belief, to imagine that this must come in the form of some *extra* action that the idea of willing can invite. Even though I just *contrasted* Descartes' picture of the will (compatible with the dictates of rationality) with Williams's (lawless will), I think it is likely the influence of this Cartesian model of control that leads Williams to discuss these issues about belief in terms of our capacity for willing belief in the first place. For it is really seems very strange to deny that one can decide what to believe because one is moved to believe by evidence or by veridical perceptions. To get to this thought, one must think of a capacity for decision in terms of some *extra* act that occurs. The picture of control I began to sketch in chapter 1 and in section 1 above does not need to posit an extra step of some sort, an intervening cause, because the notion of control sketched there

is a picture of *normative* control. The contrast is between an agency that realizes the epistemic aim of having true beliefs in virtue of its reasons-responsiveness versus an agency that would require the metaphysical feat of believing anything at all in order to be properly said to have epistemic control. Williams construes doxastic involuntarism$_2$ as an issue concerning a *particular type* of *causal* control (will as an intervening cause); we construe doxastic involuntarism$_2$ as an issue of *normative* control.

If we did have to constantly exercise a kind of extra act of willing to have control, regardless of what we say about the most conscious, deliberative cases, it would immediately seem that we lack control over (or responsibility for) the vast majority of our beliefs. And it might seem that the alternative picture of epistemic control and responsibility we sketched in section one also requires this kind of extra act, or constant intervention. There we said that we can see ourselves as agents with respect to our beliefs, as not merely saddled with them, insofar as we see our beliefs as the exercise of our rational agency, and in particular, as responsive to our rational assessments and deliberations. But surely conscious deliberation and monitoring does *not* go on with respect to most of our beliefs, most of the time. Does this mean that we are similarly led to the conclusion that we aren't therefore in control or responsible for them? The idea of virtual control[16] is of help in answering this question. Rational control is in part

[16] This is Philip Pettit's term, see *A Theory of Freedom: From the Psychology to the Politics of Agency*, Oxford University Press, 2001, p. 39. Sosa gives a similar view: "A reason-endowed being automatically monitors his background information and his sensory input for contrary evidence and automatically opts for the most coherent hypothesis even when he responds most directly to sensory stimuli. For even when response to stimuli is most direct, *if* one were also to hear or see the signs of credible contrary testimony that would change one's response. The beliefs of a *rational* animal hence would seem never to issue from *unaided* introspection, memory, or perception. For reason is always at least a silent partner on the watch for other relevant data, a silent partner whose very *silence* is a contributing cause of the belief outcome." Sosa, "Knowledge and Intellectual Virtue," in *Knowledge in Perspective*, p. 240. So according to Sosa, then, we have a way of explaining how our beliefs are different from the epistemic wantons', even at the first-order level..

comprised of virtual control—control that would step in, so to speak, if things were to go contrary to that agent's rational lights, his or her standing beliefs (and desires). A capacity for rational guidance means that a person's most unreflective beliefs are still marked by the rationality of the believer. David Lewis explains:

> An action may be rational, and may be explained by the agent's beliefs and desires, even though that action was done by habit, and the agent gave no thought to the beliefs or desires which were his reasons for action. If that habit ever ceased to serve the agent's desires according to his beliefs, it would at once be overridden and corrected.[17]

Though this idea of virtual control may help us keep hold of the idea of rational guidance with respect to the vast majority of our beliefs the vast majority of the time when we are not in any way consciously deliberating with respect to them, it suggests a new threat. For the idea of control—by mechanisms that help to keep our beliefs on track—brought to mind by Lewis's language threatens to lose sense of agency and a basis to counter doxastic involuntarism$_2$. For although the idea of virtual control gives us a way to attribute rational control to an agent without some requirement of constant deliberative intervention, it is still possible that the agent's beliefs and desires are beliefs and desires that the agent finds reason, upon reflection, to disavow. If a person has reason to attribute a belief (or desire) to herself, that, by her *deliberative* lights, isn't rational (or valuable), and if the outcome of the deliberation is not sufficient to uproot the belief (or desire), then deliberation again looks like an idle wheel. Then, although we may have some variety of rational responsiveness, it is not the sort that could make us agents competent for instantiating the requirements of the conversational stance, nor robust enough to comprise an agency that could support doxastic voluntarism$_2$ (belief acquisition *is* up to us).

---

[17] David Lewis, *Philosophical Papers*, Volume 1 (Oxford: Oxford University Press, 1983) 39.

Bernard Williams's picture of free believing seemed to require an ability that would not be of much use to us as epistemic agents—an ability to believe irrespective of the truth. Instead, we have suggested that free believing be seen as a rational responsiveness. Wanting our beliefs to be subject to the authority of reason shouldn't mean that we cannot see ourselves as agents of our beliefs insofar as we see them as responsive to our reasons for having them. But this kind of rational responsiveness, to be an agency, has to be answerable to our deliberative beliefs. Virtual control means that we don't require direct deliberative control, all of the time. But insofar as there is the possibility that deliberation is ineffective, and we have reason to attribute to ourselves beliefs that we cannot avow, beliefs which are not responsive to our higher-order attitudes, not just irrationality, but alienation threatens. If this is true, then there is a significant sense in which agency and responsibility for belief is still out of reach. This is the manner in which, I'll argue, skepticism poses a distinct threat to our view of ourselves as free and rationally responsive believers.

3.3 Authorship

Harry Frankfurt's account of free and responsible actors will give us resources to further discuss what more we might want by way of agency. Frankfurt discusses the free and responsible person in terms of having a free will, though he does not think of the free will as an intervening cause (a discrete *act* of will*ing*) in the manner we objected to in connection with our discussion of Bernard Williams.[18] Rather, for Frankfurt, what is

[18] Or at least Frankfurt does not describe the free will *primarily* or constitutively as intervening cause. This doesn't mean that the free will is not a cause at all, presumably it *is*. If authorship is essential to agency, this may be a way of explaining what it is for an agent to be a cause. Velleman (*Practical Reflection*, [Princeton, NJ: Princeton University Press, 1989]) suggests that we might think of this, in turn, as a project of capturing a naturalistically respectable agent-causation; that is, agent causation that supervenes on

important for the free will is the idea of *identification*. A person is responsible in virtue of having the will he wants, by identifying with his motivating desires. There are well-known objections to Frankfurt's account[19] as one adequate to provide sufficient conditions for responsibility. One issue concerns what Pettit has called the problem of recursivity:

> Let my responsibility be mediated by a certain sequence of controlling factors and it seems that I must be responsible for every link in that chain. I cannot be responsible for something in virtue of the operation of a controlling factor for which I am not also responsible. Responsibility is inherently recursive in nature.[20]

The problem recursivity presents for Frankfurt's view is that there are cases where it intuitively seems that even where identification is achieved and the agent is not alienated from the content of her will, still the agent should not be considered responsible. This is because the agent is not seen as properly responsible for wanting that will (i.e., due to a morally warped and sheltered upbringing).[21] *Our* primary concern is with sketching facets of agency that are threatened by skepticism, and kept viable by some adequate solution to skepticism. So the objections to Frankfurt's account that are based on its insufficiency as an account of responsibility need not directly concern us.[22] The idea of identification can still be useful to us as a *necessary* condition on agency.

---

naturalistically respectable mental states, mental events and states of affairs (and so a notion wider than Chisholm's theory of non-reducible agent-causation).

[19] One of the best known concerns whether the accord of *second-order* desires is sufficient to achieve identification in the way Frankfurt wants, for we could be alienated from a second-order desire as well: see Gary Watson, "Free Agency," in G. Watson, ed., *Free Will* (Oxford: Oxford University Press, 1982); Susan Wolf, *Freedom Within Reason*, (New York: Oxford University Press, 1990), 23-45; Philip Pettit *A Theory of Freedom: From the Psychology to the Politics of Agency* (Oxford: Oxford University Press, 2001) 53-7; Harry Frankfurt, *Necessity, Volition and Love* (Cambridge, Cambridge University Press, 1999), 104-5.

[20] Philip Pettit, *A Theory of Freedom*, 11. See also Wolf on the requirement of autonomy, *Freedom Within Reason*, esp. p. 10.

[21] Wolf, *Freedom Within Reason* 75-6.

[22] There could be an analogous problem for the adequacy of identification as a sufficient condition of epistemic responsibility. Suppose we judge some person to go wrong in his beliefs by way of continually embracing the fallacy of affirming the antecedent. He might well, and probably would, identify the beliefs formed this way as endorsable by his best rational lights. But we might not hold him epistemically

Frankfurt uses the concept of identification in explanation of what is essential to a person's freely willing an action. He argues that even in cases where one could not have willed otherwise, identifying with one's will is what makes it one's own in the sense that is relevant to responsibility. We can use this thought in relation to our evolving conception of epistemic agency. Reflexivity gives us the ability to think about and assess our perceptions and beliefs. In order to believe by our own lights, we need to be able to believe and act in response to those assessments and the reasons they give us. We can say that in order to be agents, we need not just to be rationally responsive, but rationally responsive by our own lights. We'll call this property of epistemic agents a capacity for *authorship*. Authorship means that our justificatory reasons for belief trump explanatory ones,[23] that we believe for the *reasons* that support our beliefs and not by way of some other cause. Without authorship, we may have a type of rational agency, but not the deliberative agency that would make us appropriate subjects of the conversational stance we discussed in chapter 1. To elaborate on and clarify these thoughts, we'll need some further ideas about belief in place.

3.4 Deliberative beliefs, dispositional beliefs

We can begin with the vague idea that belief is often considered a dispositional state. One believes that p just in case one is disposed to act on the basis of p in various

---

responsible for these beliefs if we don't hold him responsible for having such poor faculties of reasoning. So authorship may not be *enough* for epistemic responsibility. On the other hand, the problem of recursivity does not seem much of a problem when we get things right, or have the ability to do so. There is less temptation to say, when we get things right by way of deliberation, that we are not fit to be responsible for believing correctly by our best deliberative lights because we are not responsible for wanting our reason to be authoritative.

[23] This terminology of justifying reasons and explanatory reasons derives from Richard Moran, *Authority and Estrangement: An Essay on Self-Knowledge* (Princeton: Princeton University Press, 2001), as does much of my discussion of the importance of deliberative agency. See n. 29 below.

situations, or draw conclusions or inferences of various sorts on the basis of *p*, assert that *p* when queried, and so on. If one believes the well is poisoned, one avoids consumption of the well-water, keeps others away, perhaps initiates some action to clean-up the well water; if one believes that one is bound to win the big lottery any day now, one might quit one's job, book a vacation, pick out some extravagant gifts. Of course any specific account of belief of this type needs to formulate this idea more carefully and perspicuously, making clear how this belief is to be discerned given the complicating factors of the *other* beliefs one has, what desires, possibilities for action, possibilities for *akrasia*, and the like—*ceteris paribus* clauses are an essential part of this type of account. But, in general, this *explanatory* notion of belief as a disposition to action, in an abstract form, can be contrasted with a different, *expressive* casting of belief more central to the picture of agency we are seeking.

To do this, let's look at an example. Jonathan Cohen gives the following dispositional analysis of belief:

> ...belief that p is a disposition, when one is attending to...p, normally to feel it true that p and false that non-p, whether or not one is willing to act speak or reason accordingly...The standard way to discover whether you yourself believe that p is by introspecting whether you are normally disposed to feel that p when you consider the issue.[24]

Cohen's account is not a dispositional account that offers some nuanced explanation of how to read beliefs off of behavior (whether one's own or someone else's, and giving some further allowances for the complicating factors that beliefs don't always manifest in behavior, let alone specific behaviors, etc.). For Cohen the relevant dispositions are dispositions to feeling-as-true or feeling-as-false. I'm not so sure that a *feeling* of the truth of *p* is an introspectively identifiable feeling, let alone a good account of what it is

[24] Jonathan Cohen, *An Essay On Belief And Acceptance* (Oxford: Clarendon Press, 1992), 4.

to believe some proposition. But I think that what Cohen might be after is something closer to our interests; an account of a first-personal *attitude* of belief that does not depend on the kind of behavior-based explanatory or attributive notion that a third-personal attribution of belief might. The reason this is in our interests—an account which, like Cohen's, looks to explain belief in some first-personal aspect and not just in a primarily functional or explanatory role—is that this will give us some resources to describe an *agency* in connection with belief and knowledge. But in service of this interest, though, an apparent problem with Cohen's account is that such a feeling-as-true doesn't yet begin to serve as any particular expression of rational agency that seems important to how *I* centrally regard *my* beliefs. It doesn't have any particular connection with *rationality*: a mere feeling-as-true doesn't necessarily connect with any *reasons* I have to have such a feeling. And it doesn't have any particular connection with an *agency*: introspecting some disposition to feel that *p* doesn't seem like anything *I* centrally have any connection to, beyond the fact that it is supposed to be something I find when I introspect. I can have an itch on my elbow or a grumbly feeling in my stomach, and recognize those sensations as *mine*. But although they are mine to experience, I don't in any way see myself as responsible for these sensations, as their author.[25] They may be as surprising to me and as little under my control as something I observe that is wholly *exterior* to me, like the ant that just crawled by underfoot. So although Cohen's dispositional account does describe belief as connected to some feeling I have, it does not in any way yet characterize my beliefs as an expression of my rational agency.

---

[25] I *might*, if I had done something particularly relevant to bringing them about—like wearing a wool sweater when I know it makes my skin crawl or eating the fried wontons from the corner take-out that *always* make my stomach grumbly—but otherwise this needn't be the case.

We saw in chapter 2 that the contextualist account's commitment to an error theory about our understanding the content of our knowledge attributions led to a problem of self-knowledge. Here we can see that characterization of this problem depends on how we characterize belief. In a high-standards context, the contextualist story goes, I am disposed to deny that I know the well is poisoned, or that there is a well, and it is true to say that I *don't* know these things. Do I still believe them? Thoughts of my *past* commitments to my *knowing* these things are *now* construed as knowing these things according to higher standards in play. So I no longer am disposed to accept such attributions of knowledge as true, and I no longer see an adequate basis for a commitment to them. Insofar as I don't, and insofar as we characterize my belief in terms of my expression of commitment about the well, for the same reason I don't think I *know* it is poisoned anymore, I don't *believe* it is, either. If we think of my belief in this case as the result of my deliberations about the well, as sensitive to my judgment about what I have reason to believe about the well, then if, in this context, I judge that I don't have sufficient reason to know the well is poisoned, there is a good reason for saying I don't then *believe* it, either. If I don't think of knowledge as contextual, if I am like any of us in the unenlightened state DeRose called 'bamboozled', I will think that I no longer believe what I did before. But if I do think that knowledge is contextual, then I should still be able to think, even in this high-standards context, that I can know them according to some *lower* standard. Either I see my knowledge attributions as having explicitly standard-articulating content, in which case my commitments will side with the standard I think relevant and adequate, or I will not see my attributions as having standard-articulating content, in which case I have a certain problem of self-knowledge. The

problem is that I cannot see my former beliefs, when I switch contexts, as rationally supported; I cannot see them as commitments expressive of my rational agency by my current lights. I can see myself as disposed to have different commitments in different contexts, but I cannot avow both commitments at once, unless I see them to be explicitly expressing different thoughts. When I think about what I believe, normally I think about *the well* and its being poisoned, I don't think about my dispositions in various contexts. So that from my point of view, the expression of my (current) deliberative judgments about the well enjoy a special claim to be the best account of what my beliefs *are*; they are the expression of my rational agency, and because of that these are the beliefs of which I can see myself as author and agent. These are not beliefs I am merely saddled with because if my judgment of the facts, my assessments were different, I would no longer have the reasons I do for the beliefs I have, and no longer the commitment to them that in part comprises their status as my beliefs.[26]

But of course, there is reason to wonder about the nature of the commitment that is supposed to stand behind a failure to know in a high-standards context. For the contextualist it is true to say that I don't know, in such a context. I have described how

---

[26] Errors and unclarities here wholly mine, but this general picture of rational agency and judgment sensitive attitudes owes to several sources: Christine Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996); Rachel Cohon, "The Roots of Reasons," *The Philosophical Review* 109 (2000): 63-85; T.M. Scanlon, *What We Owe To Each Other* (Cambridge: Harvard University Press, Belknap Press, 1998); Richard Moran *Authority and Estrangement: An Essay on Self-Knowledge* (Princeton: Princeton University Press, 2001); R. Jay Wallace, "Three Conceptions of Rational Agency," *Ethical Theory and Moral Practice* 2 (1999): 217-242. I also have not been here able to adequately attend to the important differences between and subtleties of these authors' conceptions of agency. There is certainly *much* more in these works for epistemologists to learn from. It is my hope that this dissertation project is a useful step in the right direction of bringing some of this work to the conversation in epistemology; I'm inclined to think that the work of these value theorists, who've thought more about agency than most epistemologists, is relevant to thinking about *epistemic* agency, despite the differences in subject matter (i.e. differences between belief and desire, theoretical and practical reason). In particular, I'd like to bring some of the insights of this literature to bear on epistemological discussions of deontological epistemic judgments, furthering the work begun in this chapter. It seems to me that some views here are implausible, forced by a rigid idea of what agency here must look like to find some wholly different work for these judgments to do.

my thought that I *don't* know *p* provides reason for the thought that I shouldn't believe *p*. Providing reason for *this* thought (that I shouldn't believe *p*) at the same time gives content to the thought that I *don't* so much as believe it. By my current deliberative lights, I don't have adequate reason to believe *p*. Either we should therefore say that I *don't* believe p (I don't give credence to it) or, if I still find myself thinking '*p*', there seems to be a failure of deliberative rationality, a rational *akrasia*. I 'have' a belief that is isolated from my ability to avow it, espouse it, see it as justified.

But it seems that the contextualist doesn't come to this conclusion. That we continue to act on *p* and remain committed to it is ok for the contextualist since for him not knowing is strictly a high-standards affair; belief and commitments remain intact; the *appearance* of irrationality comes from thinking my high-standards claims incompatible with my low-standards claims. The counter-claim of compatibility here is a semantic description. This contextualist 'happy face' solution to skepticism, though, seems to depend on prying the normative and prescriptive dimension of belief and knowledge from this semantic content, a problem we introduced in chapter 2. Whether or not this makes the semantics less plausible (again, see the discussion in chapter 2), the fact that this semantic solution is committed to an error theory will mean these problems connected to epistemic agency follow. And the contextualists have yet to explain how to resolve them. If I place particular importance in the reflective context in which my sometimes implicit beliefs are made conscious to me, and if I can not commit to them in such a context, but find myself otherwise acting confidently on them nonetheless, this, to me, will appear like a crisis for rational authority, a case of rational *akrasia*. It is a case where I find myself with beliefs, but insofar as those beliefs are unmoved by the outcome of my

conscious deliberations on them, there is a sense in which I can fail to identify them as mine. I fail to see myself as their author.

We can now better see how the anti-insulation requirement we set forth in chapter 1 is related to agency, and how contextualism fails to meet it. The semantic solution tells us when knowledge attributions are true, but fails to explain how we are to *inhabit* our attributions, how we are to think them, how stand towards them. Fogelin makes a nice point in discussing Lewis's contextualism when he questions how we are to come down from the contextualists' semantic ascent:

> It would be nice to see how this might be done...even if this technical project can be carried out, it is hard to see what good it would do. Lewis's problem, as he characterizes it, is that he wants to indicate something important about defeating possibilities that can be properly excluded, but to do so without calling attention to them. It is hard to see how semantic ascent even if acceptably carried out could do this.[27]

Fogelin's point has to do with how Lewis can simultaneously claim that our ordinary propositions are true, whilst claiming epistemology makes knowledge claims go false. My concern is not quite Fogelin's, but related to this problem of semantic ascent and descent. There is a broader problem here, not about how Lewis can state his theory without some kind of contradiction, but with what that theory says about how we understand our concept of knowledge or don't, what we mean by our words, or don't, and what that upshot is for our reflective epistemic agency. I argued that bamboozlement makes significant problems for agency in chapter 2, and have reengaged some of these reasons above: so bamboozled, we can seem victims of a form of rational *akrasia*. The contextualists can't just claim that de-bamboozlement would clear up any such problem for the enlightened, that seeing the context-sensitivity of our knowledge attributions

[27] Fogelin, "Contextualism and Externalism: Trading in One Form of Skepticism for Another," *Philosophical Issues* 10 (2000): 56.

means we know how to inhabit them. This is because they've tied the epistemological context so closely to reflective contexts *in general*. It is not some singularly-epistemological context or project that makes knowledge vanish or our claims go false, it is scrutiny of any kind that can accomplish this; epistemology just happens to specialize in scrutiny. So they can't just toss out the findings of an epistemological context, without telling us how to avoid running into problems with the erosive scrutiny of reflective contexts in general. If we toss out all such reflective contexts just because they are erosive, we will not sustain our self-conception as competent authors and agents of our beliefs. And, as I'll argue in chapter 4, the very intuitiveness of the skeptical problem the contextualists discuss gets its life from the tight connection contextualists find between epistemology and reflective contexts more generally. So we can't just write off epistemological contexts as troublemakers if we want to retain the rational authority of reflective contexts.

And further, we don't want to say that de-bamboozlement means we know how (or better: when) to inhabit our knowledge claims just because seeing them as explicitly tied to context-relative standards means we could simply inhabit the claims when we endorse the standards at hand. Besides the fact that this makes our concept of knowledge less than fully useful (insofar as we can get 'stuck' with contextual standards that seem irrelevant to our interests; recall chapter 2 section 8[28]), it also concedes to the contextualist the priority of the descriptive over the prescriptive when it comes to the

[28] And, we are still stuck with thoughts like: the Moorean argument is sound insofar as we don't think about it; when we think about it, it is no longer sound because the standards for 'knowledge' automatically rise by virtue of our consideration of BIV possibilities. The dis-utility of such an imprisoning concept is a bit astounding. If the contextualist semantics is right, then as DeRose argues, we may indeed not "always know what we mean, but mean it nonetheless"; but equally, we may simply not be able to say or mean what we *do* mean to say.

essential content or value of our concept of knowledge. This, I will argue in chapter 4, is something we needn't, and shouldn't do.

We've seen some of the problems contextualism has with according authority to deliberative agency as a result of its insulationist semantics that separate what is true to say of knowledge in more stringent reflective contexts from what is true to say in more lenient practical contexts. And we can see, in general, how an insulationism of this sort is likely to be problematic for agency. Separating deliberative contexts from practical contexts threatens the authority of our deliberative reasons. But the authority of our deliberative reasons is what makes us epistemic agents. Without such authority for the deliberative over the explanatory, we fail to see ourselves as authors of beliefs rather than passively subject to them.

3.5 Responsibility revisited

Our capacity to assess our own beliefs and those of others, to share those assessments and be motivated by them plays an important role in our own self-conception as agents. It is a picture of acting in light of reasons we have, in believing and acting for reasons. But it is not a picture of being moved *by* these reasons without room for any agency at all. This difference is one between acting by way of our capacity for reason rather than being acted upon by reasons.

Such a capacity is also not *just* important to our (self-)*conception* as agents, it is not just a *picture* of agency. This capacity is constitutive of an agency. As such, it explains how we can be responsible for our beliefs. We have a part in making them what they are. Our epistemic wanton, who has no second order attitudes towards her beliefs, can never be motivated, on reflection, by reflection, to further investigate or to revoke or revise

some belief. We are different. Reflection can lead us to stand differently towards some belief—hold it in suspicion, no longer reason to other conclusions on its basis, refuse to act on its basis. Since we can deliberatively form beliefs, we see ourselves, and others, as reasoning sources of belief. Since we hold our beliefs to be subject to reason, we must consider all our beliefs as ultimately, defeasible. So we don't just look and see what we believe by consulting or inventorying a *fixed* stock of beliefs; we also continually determine what it is that we do believe. Our beliefs are what Scanlon has described as 'judgment-sensitive' attitudes; attitudes that are governed by the subject's judgment of the relevant reasons justifying that attitude. This is what we meant earlier when we said that for agents, justificatory reasons have authority over explanatory reasons (insofar as the agent is rational and believes freely). A person is the author of his beliefs when he believes for reasons and not some other cause. Scanlon explains:

> These are attitudes that an ideally rational person would come to have whenever that person judged there to be sufficient reasons for them and that would, in an ideally rational person, 'extinguish' when that person judged them not to be supported by reason of the appropriate kind.[29]

The fact that my beliefs are judgment-sensitive means that I see my beliefs as expressions of my rational agency, since if I judged reasons for belief differently, I would believe differently. This is how I can be the author of, and responsible for beliefs, how they can be 'up to me': it is not merely that they are caused by reasons, but that they are caused by reasons *I* have. I believe for reasons I have when my beliefs are judgment-sensitive. I determine what my beliefs are by way of my capacity to judge and assess, and to be motivated to believe or not believe on this basis.

If this is right, then we can see how skepticism is a problem not just of knowledge, but also one of epistemic agency. And we can see how an insulationist response like

---

[29] Scanlon, *What We Owe*, 20.

contextualism will not help. An insulationist response recovers knowledge, but not deliberative agency. The self-regarding attitude that 'I don't know *p*' is a judgment about the sufficiency of reasons I have for believing *p*. If I am to regard myself as a self-determining rational agent, my beliefs must be responsive to my best rational lights. This means if I judge myself not to know *p*, I don't fully authorize or recommend *p*, to myself or to others. I detach the judgment that *p* from beliefs that serve as a basis for other conclusions and actions. If nothing of this sort follows from the negative assessment of my reasons for holding *p*, it is hard to see how I can see my deliberative agency connected to my beliefs or action. Insofar as I continue to act on *p*, I don't see my deliberative rationality at work. When skeptical arguments threaten to undermine the adequacy of reasons for attributing knowledge but prove ineffective in motivating any lasting conviction, deliberative agency is threatened along with knowledge itself.

In this chapter, we also considered an argument of Bernard Williams that is representative of a view held by many epistemologists that belief is not something 'up to us'. Many epistemologists have argued from this thought to the idea that we are not properly subject to judgments of responsibility for belief. If we accept this line of argument, we have additional reason to be suspicious of the potential for sustaining our view of persons as epistemic agents and authors, beyond the reasons provided by the skeptic's arguments. So much the worse, then, for this picture of agency? Before we aver that conclusion, let's return to focus more directly on the skeptical arguments in chapter 4.

Chapter 4

A Normative Problem, A Normative Solution

In this chapter, we return to focus more directly on skepticism and responses to skepticism. I first consider an objection to contextualism that we did not focus on in chapter 2; the objection comes from critics who allege that the type of skeptical problem contextualism addresses is not a significant problem to begin with. Our casting of skepticism as a problem for the efficacy of deliberative agency in chapter 3 provides some resources to defend the contextualists on this point. This discussion of the sort of problem or problems skepticism presents reintroduces some general questions about methodology; I next explore these meta-epistemological issues, comparing different methodologies for responding to skepticism. The skeptic's generalizing strategy, in particular, is met by offering interpretations of our concept of knowledge that compete with the skeptic's view. Fallibilist, externalist interpretations of knowledge have been particularly successful here. But they fail to satisfy internalists and other critics who complain that they don't really address the problem of interest, but in a manner similar to Quinean naturalized epistemology, rather change it or leave crucial problems unaddressed. A central aspect of this complaint is the externalist failure to adequately address lingering problems of agency that the skeptic leaves us with. I argue that what is needed is a way to sustain an adequate picture of agency compatible with fallibilist accounts of knowledge. I argue that the methodological basis for pursuing this goal is best met by meeting skepticism with a normative argument. The aim of this chapter is to

explain what this means, and why I think this is a promising direction for epistemology to pursue.

4.1 Knowledge-based skepticisms

Other critics of contextualism have made the point that the view fails to be of much use as a response to skepticism if it only counters what they call 'high-standards' skepticism.[1] High-standards skepticism is a form of skepticism that allows for the possibility of a 'Russellian retreat'[2]: that is, while it impugns the idea that we have knowledge, it does not necessarily impugn the idea that we still have some positive epistemic status short of knowledge—justification, or reasonable or rational belief. A 'high standards' skepticism arrives at knowledge-skepticism by way of the thought that what we might ordinarily think of as knowledge fails to rise to standards of some 'high' sort argued to be necessary. We can derive different forms of high-standards skepticism depending on the sense in which standards are heightened. Michael Williams distinguishes two forms of high-standards skepticism based on a requirement that knowledge have justification of a Gettier-proof sort: *indefeasibility skepticism* says that knowledge requires justification that cannot be undermined by further evidence; *certainty skepticism* requires justification to "rest on evidence that excludes every logically possible 'defeater'".[3] We might also, in Ungerian fashion, consider a sort of certainty skepticism that depends on a requirement for the subject's psychological conviction of the highest order. We could require indubitability or incorrigibility. Dropping the idea that knowledge requires personal

---

[1] See Michael Williams, "Contextualism, Externalism and Epistemic Standards," *Philosophical Studies* 103: 1-23, 2001 and Hilary Kornblith, "The Contextualist Evasion of Epistemology," *Philosophical Issues* 10, 2000, 24-32.
[2] Bertrand Russell holds this type of position in *The Problems Of Philosophy*, 1912.
[3] Williams "Contextualism, Externalism and Epistemic Standards," 6.

justification, we could derive a high-standards skepticism from insisting that a candidate proposition for knowledge be demonstrable or infallible (necessarily true) or self-evident.

Though many of these forms of skepticism have been philosophical concerns at one time or another, most of these high-standards skepticisms are of lesser interest today. As Williams points out, the more restricted a definition of knowledge is, the less it seems to unsettle if it is shown our beliefs generally lack such status. If skepticism is true only in the sense that knowledge attributions are false when we raise standards to their limits, why should we care? This kind of skepticism seems to reduce to a tired trick of insisting on certain or infallible knowledge. It isn't a very upsetting concession to admit we don't have *that*. Stroud's doctor analogy again comes to mind: the truth of 'There are no doctors in Manhattan' shouldn't be very surprising or crushing if we employ 'doctor' to mean 'someone who can cure all conceivable illnesses in two minutes or less'. We already knew *that*. If this is the nature of the skeptical threat, it seems to have been diagnosed and disarmed long prior to contextualism's arrival on the scene. Fallibilism[4] is the order of the day.

Hilary Kornblith suggests that a high-standards skepticism is a "deeply deviant view about the nature of knowledge, or at least, about how the term 'knowledge' should be used." But just as we might concede to Stroud's medical skeptic his deviant use of 'doctor' and then readily agree there are no such 'doctors' in Manhattan, Kornblith suggests we could give the high-standards skeptic his deviant use of 'knowledge', and then, with little alarm, concede that we lack such 'knowledge'. Kornblith says dealing with such a skeptic is:

---

[4] The idea that a person can know based on defeasible evidence, evidence metaphysically compatible with falsity of *p*.

> like dealing with the Vermonter who insists that he won't say that it is cold outside unless it is at least 25 degrees below zero Fahrenheit. If he recognizes that there are important distinctions to be made in temperatures above minus 25, and that these distinctions have a bearing on how one should interact with the world, then the only difference between him and us is a bit of charming linguistic eccentricity. It would, however, be altogether different if this Vermonter thought that temperatures above minus 25 were all on a par, and had no differential physical effects. This latter character is more than a linguistic eccentric; he is making substantive claims about the world, claims which would be tremendously important if they were true.[5]

It may even be that there are uses of 'know' that stem from a high-standards infallibilist construal; perhaps the saying "You never know, *x* just might happen" stems from ascribing such an infallibilist meaning to 'know' and then denying that a person ever really is infallible in this sense. So when Kornblith says that the high-standards skeptic's use of knowledge is 'deeply deviant', it is best not to take this to mean that there is no basis for a claim that 'knowledge' could really be used in this sense, or has ever been used this way. Rather, we should understand Kornblith to be claiming that an *exclusively* infallibilist, high-standards reading is deviant; infallibilism is deviant when understood as a general requirement for the applicability of our concept. And the difference between Kornblith and the theorist who would allow high-standards is not, though, merely an argument about how we ordinarily do use 'know', but also how our concept should be understood. More on this below. But first let's look at Kornblith's main claim about the significance of a high-standards skepticism.

The analogy of the linguistically eccentric Vermonter is meant to bring out the idea that what is threatening about skepticism is more profound than a linguistic quibble over the propriety of bestowing some well-founded true belief with the honorific 'knowledge'. Kornblith calls 'full-blooded' skepticism, Williams 'radical' skepticism, a more

[5] Kornblith, "The Contextualist Evasion of Epistemology," 26.

threatening variety that doesn't depend on how high we place the bar for 'knowledge', but instead on the idea that we fail *wherever* that bar is placed.[6] We don't so much as have any good reasons for believing what we take ourselves to know. The distinctions between what we take ourselves to know and everything else are unfounded; it's 'all on a par'. We have no better reason to believe that the World Trade Center Towers collapsed than we have to believe that they still stand; we have no better reason to think there has recently been bombing in Afghanistan than we have to think there's recently been bombing in our own living room or kitchen. After all, the demon could choose to pull down the Petronas Towers but trick me into thinking that instead the World Trade Center Towers have collapsed, and if I were a brain-in-a-vat there might be bombing going on in my vat-room rather than in Afghanistan. The Cartesian skeptic is a full-blooded skeptic because she argues that none of our beliefs are justified *at all*. This skeptic is radical because she finds no better justification for the belief that al-Qaida was behind the World Trade Center attacks than the belief that the demon was behind a collapse of the Petronas Towers. She finds no difference in degree; she is like the full-blooded Vermonter who claims there is no real difference between twenty below zero and twenty above zero.

A point by Stewart Cohen is relevant here. (We've once considered it before, in our examination of the contextualists' aims in chapter 2, section 1.) Cohen describes his view of skepticism:

> ...as I see it, combating skepticism is a matter of refuting skeptical arguments. It is not a matter of somehow proving once and for all that we know things. And there should be nothing surprising about the fact that a response to one kind of skeptical argument does not

[6] Williams and Kornblith criticize contextualism, in particular, for addressing a high-standards skepticism which fails to be of much interest. The general point, however, that what is genuinely of interest about skepticism is the problem of whether we have any justification at all, precedes this application; probably most interest in skepticism has in fact been interest in securing this justification. See, e.g., G. Harman, *Thought*, (Princeton University Press, Princeton, 1973): "Much of current epistemology is best seen as a response to the thesis that we never have the slightest reason to believe anything."

> apply to another kind of skeptical argument...it is certainly true that some contextualists have tended to overadvertise their view as a solution to *the* problem of skepticism. But, since there is no such thing as *the* problem, there is no such thing as *the* solution.[7]

This is certainly a reasonable view. We should assume that the critics of high-standards skepticisms can endorse it; their criticism does not stem from either the view that skepticism needs to be refuted head-on by a constructive epistemology[8] or the idea that all skepticisms should yield to a unified treatment. So their objection is not simply that high-standards skepticism is not *the real* skeptical problem, or that high-standards skepticism is not the *really interesting* skeptical problem. It is rather that high-standards skepticism *cannot* be an interesting problem (or can *no longer* be an interesting problem). The suggestion is that 'high-standards' skepticism is not really a problem at all.

The way Komblith formulates his objection, his example of the Vermonter, suggests that the reason high-standards skepticism cannot be an interesting problem is that what's interesting about skepticism is severity; if skepticism isn't severe, the issues we're dealing with are merely terminological. I think the foundation of the objection is *really*

---

[7] Stewart Cohen, "Contextualism Defended: Comments on Richard Feldman's Skeptical Problems, Contextualists Solutions", *Philosophical Studies* 103 (2001): 96.

[8] Robert Audi distinguishes between *rebutting* skepticism and *refuting* it, where rebutting skepticism "refer[s] to simply showing that one or more skeptical arguments is not sound or that a skeptical conclusion has not been established", whereas refuting requires showing some skeptical conclusion is *false* by "showing a positive result such as that there is (or at least can be) justification for beliefs about the external world." Rebuttal of a skeptical thesis "would entitle us to withhold it [we are not *otherwise* so entitled?]...refutation...would entitle us to deny it." See pp. 325-6 of "The Old Skepticism, the New Foundationalism, and Naturalized Epistemology" in *Epistemology: An Anthology*, ed. Ernest Sosa and Jaegwon Kim (Malden, Mass.: Blackwell Publishers, 2000): 325-334. James Pryor calls the 'ambitious' anti-skeptical project one of refuting the skeptic using premises the skeptic would allow, while a more 'modest' anti-skeptical project aims to establish 'to our satisfaction" that we can know ordinary claims (like 'here is a hand') to be true. See Pryor "The Skeptic and the Dogmatist," p. 517. Michael Williams distinguishes a *definitive refutation* from his favored approach of *theoretical diagnosis* in Williams *Unnatural Doubts*, pp. 31-40. Given this diversity (and of course there is more to be found elsewhere), clearly there may be some objections to Cohen's chosen characterization that "combating skepticism is a matter of refuting skeptical arguments". My expectation that Cohen's sentiments are reasonable and generally endorsable is based rather on the ground he willingly *concedes*—the thought that he doesn't offer a decisive refutation of skepticism, or answer any and all skeptical arguments. It is this concession that seems reasonable, as my advocacy (in chapter 1) of a catholic attitude towards approaches to skepticism would indicate.

that accepting fallibilism settles the questions about knowledge-based skepticism. In itself, the idea that only severity is interesting is questionable. If severity and surprisingness make for an interesting skeptical problem, so does intuitiveness. (I'll also argue that intuitiveness makes trouble for the 'merely terminological' assessment of knowledge-based skepticisms.)

The contextualists themselves give a good enough reason to be interested in high-standards skepticism by way of the intuitiveness of the skeptical paradox they present (the "Argument from Ignorance", see chapter 2, section 1.2). And it is not clear that fallibilism is itself already enough to do the job of clearing up the paradox. I'd agree with Cohen's assessment of the residual difficulty: "if we want to be fallibilists, the challenge is to make that position stick, in the face of the skeptical paradoxes."[9]

Part of the intuitiveness of the paradox the contextualists focus on stems from the fact that the reasoning it embodies can have its way with us outside a philosophical context. It is not just that skeptical possibilities the epistemologist dreams up tend to make us withdraw a claim to know, but that ordinary appeals to unconsidered possibilities can also have that effect. And in ordinary situations. We claim to know (ordinarily) despite unruled-out *everyday* possibilities, too (not just without doing away with the possibilities of evil demons and vathood), but can become uncomfortable (insecure about our justification for so doing) when attention is brought to bear on the fact that we have done so. This is what the contextualist is trying to explain; this is what Robert Fogelin describes as the tendency for 'scrutiny' to cause us to withdraw knowledge claims. Let's revisit the example of Fogelin's we looked at in chapter 1:

[9] Cohen, "How to Be a Fallibilist," 117. Though, of course given the residual problems I've found in the contextualist solutions to skeptical paradox, I can't agree with Cohen's subsequent confidence that "the theory of relative alternatives, properly construed, shows us just how to do that."

> Do I, for example, know my own name? This seems to me to be as sure a piece of knowledge as I possess. But perhaps, through a mix-up in the hospital, I am a changeling. I'm really Herbert Ortcutt, and the person who is called 'Ortcutt' is actually RJF. These things, after all, do happen. Given this possibility, do I know my own name? I'm inclined to say that I do not. Not only that, philosophical naïfs, namely those who do not see that such an admission may lead to forlorn skepticism, tend to agree...reflection on unexcluded remote or *not so remote* possibilities can lead us to think we almost never know the things we claim to know. As long as we maintain this 'intense view of things' we will be disinclined to think we know things or are justified in believing things that we normally accept without hesitation...this is all very Humean, for it suggests that the application of certain concepts depends on the non-linguistic fact that human beings lack the motivation, inclination, ability, or imagination to employ them in certain ways.[10]

Fogelin's example may be a little dramatic, but it brings out that we are vulnerable to this AI[11] line of reasoning, susceptible (but not always, and some people more than others) to this line of thought.[12] The appeal of the skeptical problem based on the AI stems from its connection to our more everyday vulnerabilities, suggesting that the phenomenon we're considering really is an aspect of our everyday concept, not just an artifact of philosophical misguidedness. Most *epistemologists* balk at Lewis's suggestion that: "it 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." While they may applaud his forthrightness here, they see where this leads, and they've been down the path. But the philosophical naïfs Fogelin speaks of are a different matter. It is all too easy to get *them* to concur with Lewis's thought. Is this really just philosophical bravura or, more cynically, philosophical bullying? Perhaps. If so, it seems connected to an interesting sort of epistemic bullying that goes on elsewhere,

---

[10] Fogelin, *Pyrrhonian Reflections*, 93-4.
[11] The "Argument from Ignorance". See § 2.1.2
[12] That this is so also reveals a further aspect of the oversimplification inherent in the merely methodological view (MMV) of skepticism I criticized in chapter 1. If scrutiny can cause us to withdraw knowledge claims based on more ordinary overlooked possibilities, then more ordinary applications of a form of reasoning the skeptic uses to his extreme conclusion does have motivational effects.

extraphilosophically, where claims to knowledge are at stake. For it does seem to happen that an otherwise unexceptionable claim *can* be undermined by a persistent interlocutory bully who would insist such a claim is unwarranted, if some alternative has not, or cannot, there and then be ruled out. The less confident or aggressive are, moreover, more vulnerable. I will suggest an example of this type of extraphilosophical instantiation of AI below. It seems to me that AI's connection to this kind of extraphilosophical phenomena, by itself, makes investigation of the allegedly 'high-standards' AI worth the price of admission.

Our chapter 3 discussion of agency gives us some resources, moreover, to say what the possible consequences are when this AI line of argument finds extraphilosophical success. If I can be intimidated into the thought that I don't really know that *p*, then the prescriptivity that normally follows from an assessment that I *do* know is lost. If I instead think that I don't really know, it becomes the case that I can no longer fully authorize or recommend *p*, to myself or to others. Suppose I am going to vote for a candidate from the Forces of Good party, based on a belief that her economic policies are better for the general welfare. An epistemic bully comes along and impugns the basis of my belief about my candidate's policies—suggesting that instead of the weight I'd given to economic reasons for government to counterbalance and fund economic externalities, I've failed to consider the possible effects of some mass unemployment he is now ranting about in great detail, or how the general welfare would surely suffer from such a policy in the event of a full-scale war which is likely to happen based on some military analysis he is now suggesting, of which I've heard nothing about. Suppose I think to myself: given these possibilities, (which in truth, I'd never really thought about before, being more

inclined to worry about people who don't have health insurance than the military plot of which he spoke), do I really know the Forces of Good have the better economic plan? I may well conclude that I do not. You may object: this example sounds more than a bit silly, the way I've put things. Why should this ranter's reasons to support the Party of Evil give me mental pause? But these things happen all the time. When they do, and I am no longer confident that I know that *p*, I may no longer act on the basis of *p*, and am certainly less likely to act confidently on the basis of *p*, or recommend *p* to others, since I can no longer fully recommend *p* to myself. That such things happen stems from vulnerabilities the AI suggests, vulnerabilities connected to aspects of our concept of knowledge that give us ample reason to pursue philosophical investigation of knowledge-based skepticisms.

The critics' suggestion—that a knowledge-based skepticism that concedes something of reasonable belief or justification has conceded all we really want from knowledge—really merits further consideration, not assumption. Kornblith says that we should see the skeptic who gets us to refuse calling our beliefs 'knowledge' like we see the Vermonter: traders in linguistic eccentricity. Given that we recognize differences yet in epistemic status between our beliefs, finding some better off than others, and the Vermonter "recognizes that there are important distinctions to be made in temperatures above minus 25," any further 'difference' over what counts as knowledge is merely terminological. But it is also the case that recognizing differences between say, infants, adults and fetuses doesn't mean that there is no significance to whether we will call a fetus a *person* or not. We accord a certain status to persons; this is why it matters whether a fetus is a person or not. Likewise, we accord a certain status to knowers; this is why it matters whether we

will call someone a knower or not. This is not merely terminological, but concerns the success conditions on this particular epistemic status.

There is something peculiar about the Russellian retreat, and maybe even inconsistent, and this perhaps for some of the reasons Kornblith gives. I give some further suggestions of my own on this below in section 3. But surely the issue here is *not* merely terminological. The philosophical discussion in the environs (about closure, transmission, contextualism, and so on) is particularly lively, with far too little philosophical consensus to warrant this kind of dismissive judgment of knowledge-based skepticisms in general.

Kornblith's idea that it *is* terminological likely derives from the thought that it is reasonable belief or justification that is regulative, action-guiding, and so the more important notion. The Vermonter recognizes the difference between twenty above zero and twenty below zero and that such distinctions "have a bearing on how one should interact with the world"; we might add: 'and it is precisely *these distinctions* that have such bearing, nothing further.' If Kornblith has this thought in mind, it is an argument that Mark Kaplan gives more explicitly in his aptly-titled "It's Not What You Know that Counts."[13]

Kaplan suggests, in particular, that the pursuit of a Gettier-proof analysis of knowledge is a useless task. Useless, not because it cannot be accomplished (though that may well be true), but more literally useless—of no possible value. A Gettier-proof analysis of knowledge wouldn't help us to an improved understanding of how to pursue inquiry. After all, the premise of the Gettier cases (and here I don't mean just the cases Gettier himself offered as counterexamples to the traditional justified-true-belief analysis

[13]"It's Not What You Know that Counts," *Journal of Philosophy* 82 (1985): 350-63

of knowledge, but additionally the further family of counterexamples inspired by Gettier's 1963 article) is that I could responsibly come to my true belief that *p* whilst I still fail to *know p*. My failure to know in Gettier cases is no failure that I am to be faulted for—I am not to blame for my belief in *p* and not to be faulted for thinking I have good enough reason to believe *p*. The evidence within *my* purview justifiably supports *p* and so I would not be at fault for believing my belief in *p* to be justified. The Gettier-type facts that make it that I fail to *know* that *p* are beyond my purview. So in my own case, knowledge of what additionally needs to be the case for my justified true belief to constitute knowledge, knowledge of what de-Gettierizing attribute[14] of *p* needs to be in place in order that my true belief *is* knowledge, would not itself affect my inquiries. Knowledge of some such correct analysis could not help me *better* pursue inquiry, since we've already conceded that I've done well enough here. From the first-personal point of view, pursuing well-justified beliefs and pursuing knowledge come to the same thing.

This distinction between beliefs that are justified and beliefs that constitute knowledge gets its grip from the third-personal point of view. Surely I am saying something is lacking in your beliefs, when I fail to accord them the status of knowledge, even while I admit them to be true and justified.[15] *I* see that you are Gettiered, and so judge that you don't *really* know that *p*. Kaplan argues that even from the third-personal point of view, having a Gettier-proof analysis of knowledge to hand would not be very helpful. It isn't necessary because, according to Kaplan, the failures in Gettier cases are

---

[14] suggestions here have included versions of reliability, causal connection, indefeasibility, no false evidence or false premises.

[15] Here I retain the Gettier conception of a justified belief as a belief that is justifiably believed. It is better to more clearly mark the distinction between justification in believing (the deontological notion that assesses the believer's epistemic performance with an eye to the evidence available from the subject's point of view) from the general notion of justified belief, which needn't be so tied to what is available from the subject's point of view, but can be used to mark a (more purely objective) relation between belief and the worldly facts.

not the kind of failures that are appropriate to assessing the subject's epistemic performance, her performance as an inquirer; we have already admitted that a Gettier subject's beliefs are justified and that Gettier-related failings are beyond her purview. Moreover, we don't need a Gettier-proof analysis to assess these failings, which are fairly obvious from the third-personal point of view. We don't need such an analysis to tell us, for example, that we shouldn't accept arguments that reason from false premises. Kaplan argues:

> What is needed...is...a ground for thinking that Gettier counterexamples are of sufficient importance, indeed that *knowledge attribution* is of sufficient importance, to require an adequate theory of rational inquiry to offer an account of knowledge attribution that successfully deals with Gettier's counterexamples. But it is just such a ground for which we have searched in vain.[16]

Kaplan's ultimate conclusion is that "the moral to be drawn form Gettier's counterexamples is that what you know doesn't matter."[17]

So there is a convergence between Kaplan's line of thought and Williams and Kornblith's complaints about knowledge-based skepticisms. Neither thinks that *knowledge* in particular is all that important. Especially if what distinguishes knowledge from justification is some high-standard affair—whether that pressure for high-standards derives from skeptical arguments or from Gettier considerations. For we might understand Kaplan to be arguing that if Gettier has shown knowledge to require something *more* than justified true belief, whatever more is required by knowledge is unimportant because it is unimportant to inquiry. Kaplan thinks the high standard for knowledge that being Gettier-free requires is irrelevant to proper pursuit of inquiry, and therefore unimportant. Kornblith and Williams argue that if our skeptical worries revolve

[16] Kaplan, "It's Not What You Know," 360. Italics are Kaplan's.
[17] Kaplan, Ibid., 350.

around some high-standard concession of knowledge, such skeptical worries are of little interest or import because we don't really need *that* kind of knowledge.

4.2 But knowledge *is* action-guiding

My line of response to both Kaplan and the critics of knowledge-based skepticism will be that the action-guiding nature of knowledge is a central fact we must keep in mind, and more fully respect. This response, in the first instance, most obviously applies to Kaplan's arguments, but I will also argue on this basis against Kornblith and Williams's dismissive attitude toward knowledge-based skepticism.

First, a somewhat tangential point: this argument of Kaplan's about the general worth of a Gettier-proof analysis of knowledge overlooks the fact that, despite the sometime tedium of the Gettier-inspired literature, we've learned a lot along the way. We've learned about the myriad ways in which a true belief based on good evidence can fail to be fully certifiable, fail to be knowledge.[18] But my central dissatisfaction with

[18] These are ways in which a belief can still be only accidentally true, even though justified. This could be because its truth is not explained by, not adequately related to the (good) reasons for which it is believed, or because it is only accidentally true that belief is believed (for good reasons). This second possibility is that of misleading evidence one is lucky not to have come across (see Keith Lehrer and Thomas D. Paxson, Jr., "Knowledge: Undefeated Justified True Belief," *Journal of Philosophy* 66 [1969]: 225-37 and Gilbert Harman *Thought* [Princeton: Princeton University Press, 1973]). The first category has been subdivided in many different ways, by different authors; Fogelin talks about "weakening inference" and "epistemic-luck" examples (see *Pyrrhonian Reflections*, p. 23 ff.), Stephen Hetherington about "strange occurrences" that subdivide into "helpful" and "dangerous" possibilities (see *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology* [Oxford: Oxford University Press, 2002]: 73-5). I will further discuss some of these sorts of examples and possibilities below. My point here is that the Gettier literature helps us understand what we think it is *for* reasons to be adequately related to the truth of beliefs by classing cases where we judge there to be some undermining insufficiency. It is true that since a Gettier case is one where the subject isn't aware of the Gettier-facts (that function both to make her belief luckily true, but also in a sense, to make it unluckily almost-false), she wouldn't inquire differently in the very case at hand just by knowing some account of Gettier-type hazards that could befall an epistemic subject. At the same time, I'm not convinced that Kaplan has shown that a generalized understanding, or a sharpened understanding, of epistemic ideals and standards can't have, in any sense, an impact on the conduct of inquiry.

Kaplan's view, though, is his willingness to dispense with knowledge.[19] Kaplan argues that there are no good grounds for the importance of *knowledge attributions* in particular when it comes to thinking about good pursuit of inquiry. Kaplan argues, in effect, that *knowledge* shouldn't be our fundamental prescriptive notion, because how we evaluate ourselves and others has to do with epistemic performance, and Gettier problems show that we can't be responsible to all that knowledge in particular asks of us. Obtaining *knowledge* isn't necessarily within our control, but (personal) justification is, so we shouldn't worry about knowledge, but justification instead. But this view, as much as the contextualist view of knowledge I criticized in chapter 2, promises to divorce the prescriptive function of knowledge from whatever descriptive content it captures. Kaplan would allow epistemologists whatever best account of 'knowledge' they can come up with as descriptively adequate to our concept. He only argues for the relative unimportance of *that* concept itself. But surely this gets things backwards. The prescriptive, authorizing role that knowledge plays for us is central to our concept itself. Kaplan doesn't care much about knowledge in particular because he underestimates the prescriptive work that knowledge does and overlooks the way in which knowledge is, therefore, action-guiding. The fact that knowledge is centrally a positive epistemic status, and *a centrally-important* epistemic status that licenses action and belief, is as germane to an adequate understanding of our concept as a coherent saving of the phenomena with regard to our intuitions about its descriptive content.

[19] We might recall, a similar willingness on the part of David Lewis (see chapter 2, section 8). It is interesting the frequency with which this view arises amongst epistemologists.

Knowledge *is* action-guiding. I may pursue justified beliefs, unable to see any 'distinguishing mark' of knowledge.[20] But if I subsequently am given reason (or, led to judge, even in the *absence* of good reason) to believe that I don't *know*, this will surely have undermining effects. Of course, as I have argued previously, typically I won't *just* now say that I don't know, but still have justified belief. My tendency to judge that I don't know means that I no longer see my belief as fully justified. But if there is something about *knowledge* that allows a particular route to belief's being undermined (or partially undermined—we might say it is at the very least *decertified*), it is worth investigating knowledge in particular (not simply justification). For insofar as such undermining of belief *and* belief-guided action operates by way of our concept of knowledge, it is in here, in this way, that *knowledge* in particular *is* action-guiding, is regulative. The fully positive epistemic status conferred upon knowing, and the prescriptions (to oneself and to others) licensed by this status mean that it does matter where and when we decided to accede to such a description.

4.3 Decertification

We should worry about a tendency to demur from knowledge on the basis of high standards, because the issues here are not *just* descriptive (which *would*, if the case, rather suggest that if you concede degrees of justification, call it what you like). If high-standards skepticism finds purchase (as the AI suggests), the prescriptive weight of 'knowledge' goes by the wayside, even if there is a Russellian attempt to salvage some lesser degree of reasonable belief. We might say that knowledge is so important (and not

[20] Kaplan calls attention to this Cartesian conception of knowledge that holds knowledge to have distinguishing marks and features such that one *could* differentiate mere opinion from knowledge. "It's Not What You Know," 361.

just reasonable belief), because knowledge is fully certifying, and a denial of knowledge is decertifying . Knowledge is the terrain we want to occupy because knowledge is fully authoritative. I am arguing that this is centrally important to our use of the concept, and that this use be a *starting* point for epistemological discussions of knowledge. I'll return to this issue below.

4.4 Knowledge-based skepticism, again

But Kornblith and Williams needn't deny that knowledge itself is a centrally-important prescriptive notion. They just don't think *high-standards* knowledge need be, because we shouldn't allow a high-standard interpretation of, or requirement on, knowledge. We recall that Kornblith says that allowing a high-standard for knowledge produces a "deeply deviant view about the nature of knowledge". Williams says that "it is not clear that knowledge-specific scepticism amounts to more than fallibilism, which is less a problem than a rationally anti-dogmatic outlook...scepticism is clearly a problem only if it is radical."[21] So the common thought is: if we reject infallibilism (and we should), we shouldn't be interested in knowledge-based skepticisms; we simply accept that infallibilism is a 'deeply deviant view' and move on. But the claim that infallibilism is deviant is too quick and easy, given that our intuitions about knowledge in general are so difficult to sort out. I do think we should want to move beyond what infallibilist intuitions we do have, but seeing *how* to do this may be no easy matter. A claim that our concept of knowledge simply is *not* infallibilist is not only questionable, but also, I think, ultimately less than *helpful* in getting *beyond* high-standard intuitions. I will go on to suggest that our best bet is to be upfront about arguing *for* fallibilism on a *normative*

[21] Williams "Contextualism, Externalism, and Epistemic Standards," 6.

basis. But before moving on to this argument, I'll make one more attempt to argue that our concept of knowledge is significantly vulnerable to the pressures of high-standards. I think (some) Gettier cases, accepted by virtually all epistemologists to show that we lack knowledge when Gettiered, seem to also suggest that knowledge *is* significantly a high-standards affair. Our susceptibility to AI is similar to our susceptibility to accept Gettier intuitions; there is a pressure towards high standards at work with both.

In Ginet's Gettiered barn case[22], the subject has lots of good evidence that the barn she points to *is* a barn—it looks just like a barn, it's in the countryside near hay and horses where it makes sense to have a barn, there is adjacent farmland, etc. And, in fact, the subject doesn't just have some good evidence, she has a *true belief* that she's formed on the basis of this evidence; the structure she points to *really is* a barn. But her belief that the barn is a barn is not *knowledge* because, as the Gettier-style story has it, unbeknownst to our benighted subject, there are lots of barn facades, faux-barns, in the vicinity. She only just *happened* to point to a real barn, but it might just as well have been a mere facade.

Our Gettiered subject is not at fault for believing she's pointing out a barn; in fact she is indeed pointing out a barn. We won't want to say she *knows* it's a barn, though, because we think her belief, though true, is too *precarious* to be knowledge. There's just too much lucky compliance on behalf of the world serving to make this belief true here; she might well have had a false belief instead. There's too much fallibility in this belief's basis, given all those fake barns hanging about, for us to want to grant that that our subject knows she's pointing to a barn.

---

[22]This example is due to Carl Ginet, but popularized in Alvin Goldman's "A Causal Theory of Knowing," *Journal of Philosophy* 64 (1967): 355-72.

So far, we don't have much footing to say that we have some high standard at work here. We don't fault our subject for believing on inadequate evidence, demanding some unreasonably high standard for reasonable belief. Rather, our subject's epistemic position was poor because of all those fake barns. A world in which she would falsely believe she's pointing to a barn was, in this case, very close by indeed. But some other Gettier cases seem to more clearly pull on high-standard intuitions. Like the fake barn case, the subject here is also Gettiered by way of information that we are aware of, but that the subject is not aware of.[23] But if we judge that some Gettiered subject doesn't *know* just because it seems his belief is, from our enhanced epistemic position, formed on an inadequate evidential basis, even when epistemic danger is *really* no closer at hand, this does seem like we begin to impose a high-standard requirement on knowledge. In such cases our denial of knowledge seems directly related to the thought that the subject is not in an adequate epistemic position simply because there is some better one to be had. If this thought is behind a denial of knowledge in such cases, the reasoning at work is quite similar to the reasoning behind the AI, and a further example of how we can be pulled in by such reasoning in thinking about knowledge.

Stewart Cohen gives an example in presenting his case for contextualism that we could easily Gettierize. He describes a situation in which a couple are wondering whether a particular airplane flight makes a stop in Chicago. The wondering pair overhear one passenger ask another this very question, and this second passenger, Smith, responds to the first after checking his printed itinerary: "Yes, I do know—it does indeed stop over in Chicago." Our wondering pair are not satisfied with what they've overheard;

[23] Fogelin gives a very good description of this 'informational mismatch' behind Gettier examples. See Fogelin *Pyrrhonian Reflections*, Part One.

they have a very important drug connection to make in Chicago who won't be pleased if they don't show up.[24] They worry that the itinerary could be out of date or have some misprint, and don't think that passenger Smith really *knows* the flight will layover in Chicago; they want to get this information from the gate attendant.

Cohen uses this case to argue against the plausibility of the singular appropriateness of either Smith or our couple's standards; he argues that context determines which standard applies and determines the truth conditions for knowledge here. Anti-contextualists have responded by claiming that the lower standard really *is* the appropriate one. They claim Cohen produces intuitions in support of his contextualist analysis only by underdescribing the case.[25] Their suggestion is that if we fill in that externalist conditions for knowledge are satisfied—e.g. we do know that the itinerary *is* reliable, and that this is generally an entirely reliable method for knowing about flight layovers—then we'd no longer be tempted to say that Smith doesn't know, just on the basis of some remote possibility of misprint that hasn't obtained in this case.

Now we *could* instead Gettier the airport case producing a stronger tendency to say that Smith *doesn't* know. We might suppose that Smith is actually reading an older itinerary that his mischievous son swapped with his current (similar) flight information, or that there is a prank Smith knows nothing about where lovers of the windy city have made it so that every itinerary that day shows every flight stopping in Chicago. Gettiering this case, done properly, means that Smith is correct that the plane does stop in Chicago, Smith is justified in his belief that it does stop, but the outside information *we*

---

[24] Well, Cohen says Mary and John have "a very important business contact"; but this interpretation gives the needed urgency to their concerns. See Cohen "Contextualism, Skepticism, and the Structure of Reasons," 58.

[25] I've heard this suggestion several times in conversation; one place it can be found in print is in Matthew McGrath, "Resisting Contextualism," unpublished manuscript.

have about how his belief *could* have failed to reflect the truth makes his belief too precarious for us to grant that it is knowledge.

But again, just like other Gettier cases, the fact that our subject is justified in his belief means that an awful lot of his evidence for believing *is* in order. We focus on what Smith doesn't know and what makes his belief precarious and say that he doesn't really know. But this analysis of Gettier cases, while connecting with real intuitions most people have about these types of cases, does tend toward a high-standards view of knowledge. That this is so can be shown by a further twist in the story.

Suppose that the windy-city pranksters had actually largely failed in their computer-hacking mission to change all airline itineraries to read a stopover in Chicago. It had worked at first, but within minutes the problem was discovered and their plot foiled; only a minuscule handful of unlucky travelers had hoax itineraries.[26] But the overly vigilant and somewhat confused gate agents nonetheless, hearing of the hoax, decided to make repeated announcements over the airport loudspeakers advising caution about computer-generated misprinted itineraries. Smith, wearing headphones, is oblivious to these goings-on. Checking his itinerary, seeing that it lists a layover in Chicago, does he know that his plane will stop in Chicago? If we focus on the pandemonium in the airport, the number of other people worried about their itinerary's accuracy, those who go on to check it with the gate agents, and Smith's obliviousness, we are likely to say he does not know. But given the actually very negligible chance of itinerary error we have stipulated, the thought that Smith doesn't know seems to depend on a high-standard thought about

[26] And we could further remove question of this tiny chance making itinerary checking an unreliable route to knowledge by supposing that Smith's travel agent was tipped off to the prank and double checked all his clients' itineraries himself. Given this uber-travel agent, we can suppose that Smith has *very* reliable itineraries.

knowledge. All this information about itineraries that we have been thinking about makes it seem like Smith is not in a good enough position to know just because there is, informationally, a better position to be had: ours. We can notice that in this airport case, unlike in the barn case, actual epistemic danger is no closer for Smith than usual. So it seems that insofar as we are tempted to deny Smith knowledge, it is because all this attention to mere possibilities of epistemic danger has served to raise evidential standards for attributing knowledge.

In general, when we judge that some Gettier subject lacks knowledge, we judge that his belief is only luckily true. The truth of his belief is not adequately related to the cause of its being true for him to know, or his reasons aren't adequately explanatory of its being true for him to know. Saying that Smith doesn't really *know* whether his plane will stop in Chicago, we judge that the itinerary's actually being good evidence for the layover is just good luck. We are tempted to say this because it seems like Smith's epistemic position isn't really as good as it could be; he hasn't, say, gathered information from the gate agents to check his itinerary. It turns out, since his itinerary *is* reliable, that he isn't in actual epistemic danger. But it is only good luck that he didn't come across the misleading evidence that would throw the evidential relation into doubt. So it is not some doing of *his* that has kept his beliefs in good order. Our view of Smith is like Mr. Magoo—his obliviousness turns out not to matter. Smith's beliefs turn out to be safe, but no thanks to *his* doing. Focusing on our superior grounds, we won't want to recommend Smith's; we won't want to say he knows.

Similarly, instantiations of AI draw attention to the possibilities of epistemic danger for my belief. That in itself doesn't change the strength of my epistemic position or how

far away such dangerous worlds really are. But they introduce the thought that I could have better grounds than I now do, better justification than I now do for my belief. This thought (that I could have better grounds) casts aspersions on the sufficiency of my current justification even when it itself does not change my epistemic position. AI exploits these different senses of having better or worse epistemic positions to produce the thought that I don't know, since there are better epistemic positions to be had.

These remarks are only meant to attach the infallibilism that may seem to be the sole interest of AI to other difficult epistemic issues. I don't intend to offer any particular analysis of knowledge, or justification, or Gettier problems, or AI. But I think that it is easy for us to (sometimes) have high-standard intuitions about knowledge, and that seeing how to be a fallibilist is no easy matter, given these intuitions. I find it is somewhat disingenuous to simply say that infallibilism is a 'deviant understanding' of our concept of knowledge. Fortunately, there is a better way.

But first, we need to backtrack a bit.

4.5 Some methodological taking-stock

Having raised this question of what may be of interest about skeptical arguments that stem from AI, it is worth revisiting and focusing some of the methodological issues we considered in chapters 1 and 2.

4.5.1 Significance revisited

In chapter 1, we discussed the fact that certain types of arguments against skepticism seem less plausible in the wake of Stroud's arguments in *Significance.*[27] In particular, the

[27] Stroud, *The Significance of Philosophical Scepticism.*

inadequacy of confronting skeptical arguments with a dismissive boldness of Moore or Quine is more widely acknowledged. The Humean idea that skepticism just *isn't* a problem because, by nature, it *can't be* is not, itself, a philosophically rewarding response; it leaves the skeptic's theoretical conclusion intact. And anti-skeptical arguments that rely on the idea that skepticism is unintelligible or self-defeating fare poorly in meeting the explanatory standard now widely accepted. The skeptic's possibilities *seem* possible, perfectly comprehensible, and are at least as intuitive as possibilities as a theory of meaning that would explain that they are *not* in fact possibilities. The standard for adequately meeting skepticism is an explanation of not just why skepticism is not the problem it seems, but also why it seems to be the problem it is.

4.5.2 How to read a concept off of practices

Kornblith says that allowing a high-standards requirement produces a deviant understanding of knowledge. I've argued for the idea that high-standards are not *simply* a deviant understanding and that the AI is not only intuitive, but is a form of reasoning that, in the form of epistemic bullying (or even self-instigated worrying), can find purchase extraphilosophically, and lead to the withdrawal of a claim to knowledge. I've also argued that such a withdrawal has to be significant because refusing to ascribe knowledge is a form of decertification of belief. And because knowledge is certifying, it is action-guiding. When it is *not*, when we decertify a belief but still reason and act on the basis of that belief, we can only see ourselves, in this regard, as failing to live up to the demand that we be guided by our best rational lights .

But in spite of the arguments I've given for our vulnerability to high-standards reasoning (and the significance of this vulnerability itself), it *is* true that our practices of knowledge ascription, read *broadly*, support a fallibilist understanding of our concept of knowledge insofar as almost all the things we say people know they know (if they do) *without* having ruled out many alternate possibilities. If we want to preserve the truth of most of our attributions of knowledge, we won't say that infallibilism or demonstrability is a requirement (See chapter 1, section 3). But the skeptic doesn't argue with the facts about how we mostly do grant knowledge, but rather would argue that there is no *reasoned basis* for the distinctions that we draw. The skeptical strategy for a seemingly deviant reading can come via a *generalizing* move.[28] Such a strategy argues that the general features that seem to comprise the underlying rationale for an attribution don't in fact hold even in what are ordinarily thought paradigmatic instances. We looked at this move (though not explicitly by this name) in chapter 1, section 6. There we preliminarily suggested some features that the skeptic could employ to argue that we really lack knowledge in even what we thought were obvious cases—one candidate we looked at was an indistinguishability principle (you don't know p, if your evidence for p is the same as your evidence for q, some alternative to p). If we accept that this principle is a significant normative principle implicit in our understanding of 'knowledge', the skeptic can then argue that *consistent* application of this principle, consistent adherence to it, should lead us to withdraw previously accepted knowledge attributions. The skeptic

[28] This terminology is suggested by R. Jay Wallace, who calls this a 'generalizing strategy' in connection with the free will debate: *Responsibility and the Moral Sentiments* (Cambridge: Harvard University Press, 1996), 115. Susan Wolf gives a similar explanation in *Freedom Within Reason* (New York: Oxford University Press, 1990), 6 and she credits Thomas Nagel's "Moral Luck" (in his *Mortal Questions* [Cambridge: Cambridge University Press, 1979]) with suggesting this analogy between epistemological skepticism and skepticism about responsibility.

wants our reflective acceptance of some normative principle (if not an indistinguishability principle, some other undermining requirement) to change our intuitions about particular cases, to undermine our prereflective judgments about knowledge ascription. If we endorse the skeptic's undermining normative principle, but are unable to fully unseat intuitions in particular cases (or unable to sustain our reflective judgment of such cases), then it seems we fail to live up to our self-conception as rational agents, authors of belief that is responsive and responsible to norms we endorse.

This description of the skeptic's argumentative strategy suggests an obvious basis to refuse the skeptic: refuse to endorse the normative principles he suggests. We've already discussed the idea that we can't *just* baldly do this (chapter 1, section 1). The skeptic's case is too good to make brash refusal intellectually acceptable. Instead, one possibility is to offer an alternative interpretation of how to consistently construe the principles that underlie our concept of knowledge. The skeptic's normative proposals may not seem as plausible if we have to hand some other reflectively coherent understanding of that which grounds uses of our concept. This is just what many externalist accounts of knowledge promise to deliver. If we can accept that some wholly externalist account can make adequate sense of and ground the conditions of applicability of our concept of knowledge, then we won't be tempted to endorse conditions like the skeptic's indistinguishability principle.

But as successful as externalist accounts have been in reorienting much of the debate within epistemology, critics of externalism have presented some powerful arguments that our epistemic concepts just aren't wholly externalist. Laurence Bonjour has offered the best known examples in support of the idea that knowledge (and justification) requires

epistemic responsibility, not just reliable truth-tracking or the right causal connection to the facts.[29] Even more directly to the point at hand is the argument that accepting externalist interpretations of knowledge (and justification) just invites skeptical threats one 'level' higher up. If we accept externalism, we may see how we can know, despite the skeptic's arguments. But, can we know that we know? If not, how are we better off? Worrying about these questions is tantamount to accepting internalist interpretations of our concepts all over again. Otherwise, accepting externalism, we should also accept that we can know that we know, or justifiably believe that we justifiably believe. The fact that these problems seem to come up again suggests either that the externalist interpretations aren't adequate in the first place or that there is a problem that is simply being deferred by the externalist analysis. I argue below (section 4.6) that we shouldn't put our energies into a struggle over whether internalism or externalism is descriptively accurate to our pretheoretical epistemic concepts. So I accept the second alternative: there's a problem deferred by the externalist analysis. Not surprisingly, it's the problem of epistemic agency. But before we can elaborate here, I want to say a bit more about these methodological issues.

4.5.3 Should we look to our concept at all?

We've discussed Kornblith's complaint that a high-standards understanding of knowledge is a deviant understanding, and cast the skeptic and epistemologists as rival interpreters of the normative principles our concept of knowledge commits us to. But Kornblith's view, and way with the skeptic, is actually more radical than I have so far suggested. Kornblith is not actually so much concerned with our concept of knowledge.

[29] See Bonjour's 'clairvoyance' cases in *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), Chapter 3, 34-57.

He says he is rather interested in *knowledge itself*, and *not* our concept of knowledge. Our concept may well be internalist, infallibilist, incoherent, committed to skepticism, or all of the above. But we shouldn't care so much about investigating our concept, but rather knowledge itself. What we want an account of is the natural kind to which our uses of 'knowledge' refer. We shouldn't see our concept as laying out the conditions for knowledge; we should instead see our uses as reference-fixing. Our *concept* of knowledge may prove as inaccurate in determining the real extension of knowledge as a prescientific *concept* of gold is.

> Just as an account of our concept of aluminum is of little interest to chemistry, an account of our current concept of knowledge, I believe, should be of little interest to epistemology. Precisely because the concept of knowledge currently at work may be founded on a good deal of ignorance and mistake, even if there is some real understanding of the phenomenon built into it as well, we do best when we investigate the phenomenon of knowledge itself rather than our current conception of it. As with science, the point of philosophy is not to understand our concepts but to change them.[30]

4.6 A normative argument

I am sympathetic to this last, reformist thought of Kornblith's. But Kornblith means that we change our concepts by better understanding what our best science reveals; our concepts change for the better in the wake of better scientific understanding. Kornblith starts from the assumption that knowledge is a biological phenomenon. He then reasons that the best understanding of knowledge will be one grounded in the best *scientific* understanding and scientific uses that are found for it, a task Kornblith himself pursues by way of looking into how cognitive ethologists understand what knowledge is in theorizing about bird behavior. But if knowledge is importantly evaluative (as I have stressed) and moreover action-guiding insofar as it is normative for us, then we should

---

[30] Kornblith, "In Defense of a Naturalized Epistemology," in *The Blackwell Guide to Epistemology*, ed. John Greco and Ernest Sosa (Malden, Mass.: Blackwell Publishers Ltd, 1999), 166. See also "Knowledge in Humans and Other Animals" *Philosophical Perspectives* 13 (1999): 327-346.

inquire into the interests we have in the epistemic evaluations we make. And we needn't assume that our interests are exclusively biological. We assume this about birds, but we don't assume this about ourselves.

Kornblith's general reformist comment is nonetheless helpful: philosophy needn't just seek an understanding of our epistemic concepts, conceived of as static and given. The best way to approach the seemingly intractable debate between the skeptic, the externalists, and the internalists over the 'correct' reading of our concepts from our practices is not just by proceeding with examples, counterexamples, and further argument about what intuitions we *really* have about what knowledge *really* is. Kornblith is right that we needn't be *that* wedded to our concepts as we find them.

We should instead approach the debate with more explicitly normative considerations in mind, and consider any interpretative account of what our epistemic concepts are as partially normative in character—what should we want them to be? This cannot be a flatly pragmatic question. Arguments over the best philosophical understanding of our concepts generally operate via (wide) reflective equilibrium between general principles and intuitions about specific cases. Forefronting the normative means that we can give weight not just to finding general principles that square with intuitions about specific cases, but also to finding general principles we endorse as serving our epistemic ends. Sally Haslanger advocates asking such questions about the cognitive values we endorse, and calls this a form of *critical* immanent epistemology.[31]

---

[31] Haslanger, "What Knowledge Is and What It Ought to Be," 467. Haslanger is concerned in particular with simply accepting our entrenched epistemic concepts in light of the fact that we may well find them to be "sexist, androcentric, or otherwise politically problematic." She says that indeed "feminists have accumulated substantial evidence that our actual knowledge attributions and practices of authorization privilege men and help sustain sexist and racist institutions." See p. 462. So this gives her suggestion a particular practical backing and impetus.

Haslanger suggests we entertain a very broad question: "what work does it [a concept of knowledge], or (better) could it, do for us?"[32] Haslanger admits that there are core, constitutive *epistemic* values, but even here she is willing to push her normative questioning to its limits:

> It may be that some constitutive epistemic values (such as truth) can be discovered without a consideration of contextual values, while others require attention to social context. But before granting that even this small part of epistemology can proceed without attention to social and political matters, I think it is valuable to reconsider the basis for regarding truth as an epistemic value. After all, truth may be a constitutive goal of belief, but is there some reason we should see ourselves as committed to forming *beliefs* (as opposed to, say, acceptings)? Is there some value in being a *believer*?[33]

Haslanger does go on to give grounds for a positive answer here; I'm not confident I can endorse the grounds she does give. But we needn't pursue normative questions quite so far reaching as this question about belief, or agree that it is sensible to think we *can* reasonably entertain reforms in our cognitive economy that are that distant. Our concern is with the normative principles the skeptic urges as conditions on achieving knowledge properly so-called; we want to know whether we should take on the normative burdens the skeptic suggests. I can't take on the large project of exploring the far-reaching valuational questions Haslanger suggests here. But I think that nonetheless we can see our way to a superior methodological approach to answering the skeptic, without going too far off topic, and in advance of a defense of specific answers to these broad valuational questions. The skeptic defends such a high standard requirement on knowledge by the suggestion that it latches onto a constitutive value of knowledge that we won't want to give up—knowledge that is the product of rational agency. So the skeptic is himself offering a normative argument. We won't (in light of our discussions

---

[32] Haslanger, "What Knowledge Is and Ought to Be," 467.
[33] Ibid., 470.

of agency in chapter 1 and 3) want to meet the skeptic by denying the importance of epistemic agency altogether; this is not the normative strategy I'd suggest (though other epistemologists would). But since agency is territory we have already explored in chapters 1 and 3, we yet have resources to undermine this skeptical strategy. However, this is to anticipate; I still want to say a bit more by way of contrasting a normative strategy with a methodology that takes itself to be more strictly descriptive.

We have a more radical approach to the work of epistemology if we open up the questions that Haslanger suggests. Different answers to her valuational questions will potentially guide us to different assessments of regulative principles, to different sets of normative principles we endorse. If normative principles are the regulative principles which indicate how we should reason and when we should believe or believe that we know, Haslanger's valuational questions ask what master such principles themselves should serve. We don't just find out which guiding normative principles are embedded in our epistemic concepts (by reflectively matching given intuitions about cases to candidate principles). Such an analytic effort, according to Haslanger, should ultimately only serve a broader project of valuational inquiry.

But the idea that what epistemologists are doing when pursuing the analytic project is *merely* finding some descriptively adequate principle is, I think, not apt to begin with. There is good reason to think that there's some *convincing* going on alongside what is sometimes cast as a more purely descriptive project. Intuitions are being informed by, and not just reflected in, the work that epistemologists offer.[34] So part of what goes on is

[34] Stich tries to marry this idea to empirical evidence he's pursuing in ongoing research to show that reasoning from such intuitions to epistemic norms is hopelessly parochial. See Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich, "Normativity and Epistemic Intuitions," forthcoming in *Philosophical Topics*, in addition to Stich's *Fragmentation of Reason*. *Pace* Stich, the fact that our intuitions can be

not just a 'fitting' of theory to practices. The process of reflection on principles and cases can also result in one's being *convinced* of the worth of some principle (or property) via this process of reflection, and convinced of its applicability to further cases. This should sound familiar: it is the generalizing strategy I've attributed to the skeptic. I take what Stich and Haslanger would characterize as a merely analytic project in epistemology to actually be partially reformist; by 'cleaning up' our concept, our conditions of application for knowledge, theorists are engaged in a generalizing strategy of their own. Haslanger is right, at least, that many epistemologists have been not very amenable to this description of what they are up to.

I don't think there's anything wrong with proceeding in this manner. In fact, it's how we should expect to proceed if we see philosophy as sustained reflection of a piece with the concerns and attitudes of everyday life rather than some entirely separate enterprise or set of concerns. And more important to the issues at hand, it is how we should proceed if we want to effectively counter the skeptic. As I've said, the skeptic is effective insofar as he convinces that his normative principles, his high standards, are norms we *should* meet. The counterstrategy is not to explain that the skeptic offers some poor interpretation of our standing epistemic concepts, but to explain why we should *not* endorse his normative principles. An effective counterstrategy wins normatively when it successfully wins our reflective acceptance. This is why contextualism is a nonstarter in

---

shaped by epistemological theorizing isn't just a 'bad' sign of parochialism. Rather it's good news insofar as it suggests some (further) relevancy to the work philosophers do. Extended reflection on cases and principles sometimes gets us somewhere; and where it gets us is not just some (thinly disguised) road to relativism as Stich would have it. (Stich says, for example, that: "these intuitive notions of cognitive evaluation are themselves local cultural products, and there is no reason to think that they won't exhibit just as much intercultural and interpersonal variation as the cognitive processes that they evaluate. In light of this, it is hard to see why most people would *care* very much whether a system of cognitive processes falls within the extension of some ordinary notion of epistemic evaluation..." *Fragmentation of Reason*, 20.) My point obviously needs further argument. I'll have to take this up as a further project; stalking Stich here would take us too far from the issues at hand.

combating skepticism. It doesn't offer the skeptic any competition here. In ceding reflective contexts to the skeptic, the contextualist cedes the ground we want to win, if we value our reflective agency.

I mentioned above the internalist complaint that externalism 'changes the subject' in some objectionable way. If 'changing the subject' is meant to register the idea that some externalist account of an epistemic concept fails to capture our ordinary, pretheoretical concept, then, according to the line I am advocating here, this is not necessarily a serious problem at all. If our ordinary concept leaves us vulnerable to epistemic bullying, or with different standards for the anxious and the aggressive, then it may well be to the good for philosophy to play some prescriptive, and not merely descriptive role. That is, if some externalist account can explain how we need not understand our epistemic concepts in such a way as to leave us vulnerable to these everyday pitfalls, and if the concept we are left with is one that still substantively does the work we need it to do, then if this is 'changing the subject', so much the better. The question about epistemic norms should not just be: what norms are reflected in our standing concepts?, but also: what norms are in accord with what we epistemically value?, the question that Haslanger urges us to ask.

But perhaps there is some thought along these lines actually behind the internalist complaint that externalists are objectionably 'changing the subject'. Perhaps what the internalists really want to insinuate is that the epistemic properties externalists describe are inadequate to our concepts because they don't sufficiently capture what we (should) value. This might be the case, for example, if it could be argued that some externalist construal of knowledge couldn't be brought into alignment with an otherwise tenable and valuable conception of epistemic agency.

Internalists are worried about the role of epistemic agency in relation to the epistemic properties externalists describe when they ask whether we can *know* that we know, or if we justifiably believe that we justifiably believe. They are worried about our ability to know and believe for reasons, and don't find externalist accounts to mitigate such worries. Even if we *were* to accept some externalist account of knowledge, this line goes, at best we'd be left with a skeptical problem that recurs 'one level up'. We could accept that we know *given that the externalist theory is right* and *given that the world really is as we think it is*. But this, to the internalist, doesn't look like knowing, but something closer to a *conditional correctness.*[35] This problem of epistemic agency is the problem I said above was only deferred by the externalist accounts of our epistemic concepts. One way to take this up is by looking at what has been called the 'KK condition' on knowledge.

4.7. The KK condition

The KK condition for knowledge expresses the thought that in order to know that p, one must also know that one knows that p. This principle is sometimes relied on to argue for

[35]This is Stroud's phrase, but here, the 'conditions' are those of his "Understanding Human Knowledge in General," (in *Knowledge: Readings in Contemporary Epistemology*, ed. Sven Bernecker and Fred Dretske [Oxford: Oxford University Press, 2000], 307-323), rather than those of *Significance*. In *Significance*, the conditional correctness *of skepticism* is the thesis that if the philosophical question about knowledge of the external world is meaningful, the skeptic's answer to it is the right one. See 194, 214. In "Understanding," the conditions are on externalism, considered as an answer to a philosophical question about knowledge: "The scientific 'externalist' claims to have good reason to believe that his theory is true. It must be granted that if, in arriving at his theory, he did fulfill the conditions his theory says are sufficient for knowing things about the world, then if that theory is correct, he does in fact know that it is. But still, I want to say, he himself has no reason to think that he dos have good reason to think that his theory is correct. He is at best in the position of someone who has good reason to believe his theory if that theory is in fact true, but has no such reason to believe it if some other theory is true instead. He can see what he *would* have good reason to believe if the theory he believe were true, but he cannot see or understand himself as knowing or having good reason to believe what his theory says." (321) It is important to note that Stroud's argument in this paper is not against externalism per se, but against a certain interpretation of what (kind of question) externalism can be an answer to; Stroud's point is here, as in *Significance*, to question what is wanted in the philosophical project of "understanding human knowledge in general".

skepticism via meta-level doubts. Even if it is granted that I know the earth is spherical and not flat, do I really know that I know this? If I give in here, then one could argue from this basis that I therefore can't really be said to know the earth's spherical in the first place. Here, Sosa gives an intuitive presentation of the KK condition:

> Suppose that, concerning a certain subject matter, you ask yourself whether you know, and you have to answer either "Definitely not," or "Who knows?" In that case, is there not some straightforward and widely shared sense in which you do *not* really know?...It is better to believe and to act in ways that are *reflectively* right than in ways that happen to be right but unreflectively so.[36]

This second thought of Sosa's, that identifies the KK condition as a specifically *reflective* requirement, makes the KK condition an *internalist* requirement for knowledge.[37] In order to have knowledge that *p*, if we accept the KK condition, one must have some true belief (plus whatever else is necessary for that belief to be knowledge) *that one knows p*. But externalist accounts don't require that we have any belief at all, at this secondary level (in order to know *p*).[38] In any *given* case, if knowledge *does* iterate in this way, the externalist can give a consistently externalist interpretation of that fact: for example, if we accept some reliabilist account of knowledge, and grant that *S* knows that *p* because *S*'s true belief *p* is formed via some reliable method, it *might* also be true that *S* also has a true belief formed by way of some reliable method that *S*'s true belief *p* is formed via some reliable method. But the reliabilist who explains *S*'s knowledge that *p* in terms of

---

[36] Ernest Sosa, "How to Resolve the Pyrrhonian Problematic: a Lesson from Descartes," *Philosophical Studies*, 85 (1997): 229-49.

[37] Though this is not a requirement all internalists would accept. Amongst access internalists, some require access to that which justifies, while others additionally require access to the *quality* of one's justification. In the first instance access to that which secures justification is required, in the second also access to the fact of objective justification. So, for instance, some evidentialism might hold that a subject is justified in believing *p* iff his evidence for *p* supports belief in *p*. Justification is explained in terms of evidence that is available to the subject and a support relation that holds between that evidence and that proposition; to be justified the subject needn't have access to the obtaining of that support relation.

[38] Again, many *internalist* accounts wouldn't embrace this requirement, either. But if KK has any motivation at all, it would come from some internalist construal of knowledge that would anticipate that knowledge should iterate in this way.

reliability will not say some further belief needs to be in place for that reliability to obtain.

One thing that we should remark upon right away is that Sosa's intuitive presentation of the KK condition *itself* raises a second-level question. That is, we should notice a difference between:

Case1: *S* knows that *p* (let us suppose: according to some reliabilist account), but has no *further* beliefs about her knowing, or not knowing, that *p*.

Case 2: *S* knows that *p* (according to our reliabilist account), but when she thinks about whether she knows, she thinks "Definitely not" or "Who knows?"

Sosa suggests that it is difficult to think one knows if, reflecting on the matter, one is inclined to think one does *not* know, or is indifferent. But *this* seems a case in which my secondary reflection threatens to unseat my first level *belief.* In this case, it is then very easy to say why I don't know that *p*, in a way both the externalist and the internalist accept: it is because I don't even *believe* it. The way Sosa describes it, my secondary reflection concludes with the thought that I don't know, and so is decertifying, erosive in the way we've already discussed. If I am reading the paper and come to truly believe, on this basis, that George Bush choked and fainted, but then start to wonder about whether this is some fake paper, some gag, or for some other reason come to think that I either definitely *don't* know that Bush fainted or have no idea whether I know Bush fainted, it also seems that I no longer *believe* that he did. So this kind of case Sosa raises, a Case-2 situation, seems quite different, in fact, from the more basic Case-1 violation of KK. Case-2 examples seem to exemplify the kind of erosive effect of decertification that

we've already discussed in connection with AI. Sosa argues, in effect, for a weaker principle we might call the principle of self-voided knowledge:

SVK) If I think that I *don't* know that *p*, then I don't know that *p*.

This principle, if true, seems troubling for those prone to self-effacement or self-doubt. But it isn't the same as the KK condition. SVK might not affect us at all, if we are either unreflective or unperturbed when we do reflect on our beliefs. SVK kicks in when, on reflection, we have some motivating reason to doubt. If we are to object to externalist accounts on the basis of a KK requirement, we'd have to find Case-1 examples objectionable (at odds with where we think we should accord knowledge), not merely Case-2 examples. Taking Case-1 examples to be actual counterexamples to knowledge only seems tempting if we jointly endorse the following ideas: (1) Knowledge properly so called is *human* knowledge (in the sense of §3.1); (2) Human knowledge and agency are to be understood in such a way that a second-order belief is required *in each instance* of first-order knowledge. Obviously, KK requires even *more*: we don't merely have to have a second-order belief in order to have knowledge, we have to have second-order knowledge. But our chapter 3 discussion of agency should have made it clear that (1) and (2) needn't be paired. We can endorse (1) while rejecting (2). Reflective agency doesn't require constant active control (§ 3.2).

So we reject KK, and accept only SVK. I've previously argued that we can be vulnerable to self-voiding of knowledge. But here we've arrived only at a way in which reflective agency can *undermine* our epistemic status. If the externalist concept of knowledge leaves agency in question, we have in view, after looking at KK, moved no further towards knowledge that is more comfortably the product or province of some

reflective agency. Sosa declares that: "It is better to believe and to act in ways that are *reflectively* right than in ways that happen to be right but unreflectively so." This seems to express the thought that epistemic states we value are valued as products of our epistemic agency and not merely for their truth, their reliable truth, or their instrumental value in our successful functioning. But Sosa's thought is one that needs careful handling. I'm not sure that either many internalists *or* externalists handle it very well, and the skeptic is able to exploit this thought as a consequence. What is needed as counterstrategy to the skeptic, then, is not just some rival interpretation of our epistemic concepts and the normative principles that underlie them, but also a better understanding of epistemic agency. The skeptic's normative principles themselves seem valuable when taken to be necessary for the type of epistemic agency we value. Our reflective susceptibility to the skeptic's principles comes from misunderstanding the type of control necessary to be an epistemic agent.

4.8. Sensitivity and Control

The principle of sensitivity says that:

S's belief that *p* is *sensitive* if *S* would not have believed *p* if *p* had been false.[39]

If our beliefs have to be sensitive to be knowledge, without any *sotto voce* exemptions (like Lewis's; see his definition of knowledge in § 2.7) or 'nearby worlds' restrictions, our knowledge of contingent truths can't amount to much. If we are fallibilists, we allow that *S* can know based on defeasible evidence, evidence metaphysically compatible with

[39] This formulation follows DeRose, "Solving the Skeptical Problem," section 7, who takes the notion from Nozick's subjunctive conditional or "truth-tracking" account of knowledge in *Philosophical Explanations* (Cambridge: Harvard University Press, 1989). Normally, sensitivity is used in such a way that a belief is sensitive if *S* doesn't believe it in the closest (or alternatively: close *enough*) possible worlds in which it is false. Here I instead discuss an unrestricted notion of sensitivity; *absolute* sensitivity.

falsity of p. But then in some (however distant) possible world, p is false, but S still has the same evidential basis for belief in p, and S still believes that p.[40] I believe that the Hopkins philosophy department has offices in Gilman Hall, since last time I was in Baltimore, this was so. They are there, and I know them to be. But it is *possible* that the offices had moved; if they had moved I would falsely believe that they are in Gilman, though they would not be. Such a world isn't even as distant as some evil-demon or envatted-brains world. But is *it* relevant to my knowing, or not knowing, that the philosophy department offices are in Gilman Hall? This hardly seems right. My true belief shouldn't have to be sensitive to this kind of far-out possibility in order to be knowledge.

Yet the skeptic's use of an argument like AI seems to draw on sensitivity; so does an epistemic bully who relies on a more ordinary instantiation of AI to impugn knowledge. If we, on reflection, think that knowledge has to be sensitive, in some unrestricted way, we are vulnerable to the erosions of self-voided knowledge (SVK). When some specific possibility, even a not-very-credible possibility, is presented that is incompatible with p, it can seem that it must be ruled out if we are to know; (unrestricted) sensitivity can seem like a necessary requirement for knowledge, if knowledge is not to be *luckily* had. But if we assess the stringency of a sensitivity requirement on knowledge in tandem with questions about the values and purposes served by attributions of knowledge, the normative requirements the skeptic suggests as a success condition on this epistemic status no longer look very convincing. They no longer look very convincing *as long as* we *also* have a conception of agency that enables us to turn aside such too-stringent

[40] If talk of possible worlds isn't helpful here, we could instead simply think of alternative possibilities; here we'd want alternative possibilities that are compatible with the evidence, but incompatible with the truth of p.

principles. And too often epistemologists who would offer fallibilist, externalist interpretations of knowledge fall down here, suggesting, for example, that deontological interpretations of justification are misplaced. Which in turn can suggest that a deontological *component* of justification or knowledge has no place, and that we are not properly considered *responsible* knowers. I see no reason for this to be the case. A fallibilist, significantly externalist account[41] can be wholly compatible with a picture of human knowers as epistemically responsible knowers; my intent in chapter 3 was to begin to describe how this account might look. Such an account should be adequate to sustain the second-order deontological attitudes that reflect the necessary assumptions of the conversational stance (recall §1.5). The power of AI and unrestricted sensitivity derive from *both* an optional interpretation of knowledge and an optional picture of epistemic agency that supports that picture of knowledge. If we want to be contented fallibilists in the face of AI, a clearer picture of agency is the complement necessary to a fallibilist picture of knowledge. Then we'd have a reflectively coherent conception of ourselves as rationally-responsive, self-governing *human* knowers (to buttress against philosophical and extraphilosophical assaults of AI and SVK).

4.9 The value of agency: further problems and projects

The strategy I've argued is needed to meet the skeptic involves an interpretation of our epistemic concepts that is grounded in epistemic values we endorse. The skeptic's interpretation of knowledge is not adequately undermined as long as the skeptic can

[41] Here I mean a externalist account that retains some epistemic responsibility clause, but escapes principles like unrestricted sensitivity by way of externalism. (These labels are *always* vexed; no one is a complete *internalist* about knowledge since knowledge is factive. Often, 'externalism' is used to mean a view which drops epistemic responsibility altogether, which is why I speak of a 'significantly' externalist account, wanting to mark the features that allow a move away from the skeptic's unrestricted sensitivity.)

exploit principles that appear necessary to the kind of epistemic agency we value. I do think that Lewis is largely persuasive about how we have substantially fallibilist practices and substantially infallibilist reflective intuitions about knowledge. The best way against the skeptic is not, pace the contextualists, to argue for a concept of knowledge that allows such reflectively infallibilist intuitions to stand, nor is it to argue for a fallibilist concept of knowledge that would do away with a robust interpretation of epistemic agency altogether. That this conception of agency is one we should not (and likely: could not) give up requires argument beyond the considerations I've offered in this essay. One important further project that would support the argumentative line I have taken in this essay would be to further expand and argue for this conception of agency; more could be said, in particular, about the significance of the assumptions of the conversational stance, and the broad role they play in our thinking. As I indicated in passing in chapter 3, I also think that an account of our deontological talk about belief and knowledge would comprise a further part of such a study. An additional project would be to further pursue the valuational questions Haslanger proposes; this connects directly back to our first project if it assumed that further explaining the value of agency will be important here.

In advance of such projects, though, I have argued for reasons to think the route I suggest is a promising path for epistemologists who are concerned to address skepticism. Moreover, if the strategy I suggest is a sound one, we shouldn't just expect to gain some theoretical ground, but, by way of argument for epistemic concepts we value, we can expect extraphilosophical benefit from pursuing this philosophical program as well. I think such normative grounds provide the best way to confront the skeptic, and also the best way to remain, both philosophically and *extra*philosophically, a contented fallibilist.

Bibliography

Alston, William. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2 (1988): 257-99.

Audi, Robert. "The Old Skepticism, the New Foundationalism, and Naturalized Epistemology." In *Epistemology: An Anthology,* ed. Ernest Sosa and Jaegwon Kim, 325-334. Malden, Mass.: Blackwell Publishers Ltd, 2000.

Bernecker, Sven and Fred Dretske. "Scepticism: Introduction." In *Knowledge: Readings in Contemporary Epistemology*, 301-306. Oxford: Oxford University Press, 2000.

Bonjour, Laurence. *The Structure of Empirical Knowledge.* Cambridge: Harvard University Press, 1985.

Brueckner, Anthony. "The Structure of the Skeptical Argument." *Philosophy and Phenomenological Research* 54 (1994): 827-35.

Cavell, Stanley. *The Claim of Reason: Wittgenstein, Scepticism, Morality and Tragedy.* Oxford: Oxford University Press, 1979.

Cohen, Jonathan. *An Essay On Belief And Acceptance.* Oxford: Clarendon Press, 1992.

Cohen, Stewart. "Knowledge, Context, and Social Standards," *Synthese* 73 (1987): 3-26.

_____. "How to Be a Fallibilist." *Philosophical Perspectives* 2 (1988): 91-123.

_____. "Two Kinds of Skeptical Argument." *Philosophy and Phenomenological Research* 58 (1998): 143-159.

_____. "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives* 13 (1999): 57-89.

_____. "Contextualism and Skepticism." *Philosophical Issues* 10 (2000): 94-107.

_____. "Contextualism Defended: Comments on Richard Feldman's 'Skeptical Problems, Contextualists Solutions.'" *Philosophical Studies* 103 (2001): 87-98.

Cohon, Rachel. "The Roots of Reasons." *The Philosophical Review* 109 (2000): 63-85.

Dennett, Daniel. *The Intentional Stance.* Cambridge: MIT Press, 1987.

DeRose, Keith. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52 (1992): 913-929.

_____. "Solving the Skeptical Problem." *Philosophical Review* 104 (1995): 1-52.

_____. "Responding to Skepticism." In *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted Warfield, 1-24. New York: Oxford University Press, 1999.

_____. "Response to Schiffer." Manuscript on-line. Available from http://pantheon.yale.edu/~kd47/Response-to-Schiffer.htm. Accessed March 2001.

Descartes, René. *Meditations on First Philosophy.* In *Philosophical Works of Descartes*, translated by Elizabeth S. Haldane and G.R.T. Ross. Cambridge: Cambridge University Press, 1911.

Feldman, Richard. "Skeptical Problems, Contextualist Solutions." *Philosophical Studies* 103 (2001): 61-85.

Fogelin, Robert. *Pyrrhonian Reflections on Knowledge and Justification.* New York: Oxford University Press, 1994.

_____. "Contextualism and Externalism: Trading in One Form of Skepticism for Another." *Philosophical Issues* 10 (2000): 43-57.

Frankfurt, Harry. "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68 (1971): 5-20.

_____. *Necessity, Volition and Love.* Cambridge: Cambridge University Press, 1999.

Gettier, Edmund. "Is Justified True Belief Knowledge?" Reprinted in *Knowledge: Readings in Contemporary Epistemology*, ed. Sven Bernecker and Fred Dretske, 13-15. Oxford: Oxford University Press, 2000.

Gibbons, John. "Externalism and Knowledge of the Attitudes." *The Philosophical Quarterly* 51 (2001): 13-28.

Goldman, Alvin. "A Causal Theory of Knowing." *Journal of Philosophy* 64 (1967): 355-72.

Greco, John. *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry.* Cambridge: Cambridge University Press, 2000.

Grice, Paul. *Studies in the Way of Words*. Cambridge: Harvard University Press, 1989.

Harman, Gilbert. *Thought*. Princeton: Princeton University Press, 1973.

Haslanger, Sally. "What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology." *Philosophical Perspectives* 13 (1999): 458-480.

Hetherington, Stephen. *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology.* Oxford: Oxford University Press, 2002.

Hofweber, Thomas. "Contextualism and the Meaning-Intention Problem." In *Cognition, Agency and Rationality*, ed. K. Korta, E. Sosa and J. Arrozola, 93-104. Dordrecht: Kluwer, 1999. Also available on-line from http://www-personal.umich.edu/~hofweber/papers/con.pdf. Accessed January 2002.

Hume, David. *A Treatise of Human Nature*. Ed. P.H. Nidditch. Oxford: Oxford University Press, 1978.

Joyce, James M. "The Lasting Lesson Of Skepticism." Manuscript on-line. Available from http://www-personal.umich.edu/~jjoyce/383/winter00/lesson.htm. Accessed March 2001.

Kaplan, David. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals." In *Themes from Kaplan,* ed. Joseph Almog, John Perry and Howard Wettstein, 481-614. New York: Oxford University Press, 1989.

Kaplan, Mark. "It's Not What You Know that Counts." *Journal of Philosophy* 82 (1985): 350-63.

Keefe, Rosanna and Peter Smith. "Introduction: Theories of Vagueness." In *Vagueness: A Reader*, 1-57. Cambridge: MIT Press, 1997.

Kornblith, Hilary. "In Defense of a Naturalized Epistemology." In *The Blackwell Guide to Epistemology*, ed. John Greco and Ernest Sosa, 158-169. Malden, Mass.: Blackwell Publishers Ltd, 1999.

_____. "Knowledge in Humans and Other Animals." *Philosophical Perspectives* 13 (1999): 327-346.

_____. "The Contextualist Evasion of Epistemology." *Philosophical Issues* 10 (2000): 24-32.

Korsgaard, Christine. *The Sources of Normativity.* Cambridge: Cambridge University Press, 1996.

Lehrer, Keith and Thomas D. Paxson, Jr. "Knowledge: Undefeated Justified True Belief." *Journal of Philosophy* 66 ([1969): 225-37.

Lewis, David. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8 (1979): 339-359. Reprinted in *Philosophical Papers*, Volume 1, 233-250.

_____. *Philosophical Papers*, Volume 1. Oxford: Oxford University Press, 1983.

_____. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (December 1996): 549-567. Reprinted in *Knowledge: Readings in Contemporary Epistemology*, ed. Sven Bernecker and Fred Dretske, 366-384. Oxford: Oxford University Press, 2000.

Moore, G.E. "Proof of an External World." In his *Philosophical Papers*. London, 1959.

Moran, Richard. "Impersonality, Character and Moral Expressivism." *Journal of Philosophy* 90 (1993): 578-595.

_____. "Self-Knowledge: Discovery, Resolution, and Undoing." *European Journal of Philosophy* 5 (1997): 141-161.

_____. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press, 2001.

Nagel, Thomas. "Moral Luck." In *Mortal Questions*. Cambridge: Cambridge University Press, 1979.

_____. *The View from Nowhere*. New York: Oxford University Press, 1986.

Nozick, Robert. *Philosophical Explanations*. Cambridge: Harvard University Press, 1989.

Perry, John. "Indexicals and Demonstratives." In *A Companion to the Philosophy of Language*, ed. Bob Hale and Crispin Wright, 586-612. Oxford: Blackwell Publishers, 1997.

Pettit, Philip. *A Theory of Freedom: From the Psychology to the Politics of Agency.* Oxford: Oxford University Press, 2001.

Pettit, Philip and Michael Smith. "Freedom in Belief and Desire." *The Journal of Philosophy* 93 (1996): 429-449.

Plantinga, Alvin. *Warrant: The Current Debate.* New York: Oxford University Press, 1993.

Pryor, James. "The Skeptic and the Dogmatist." *Noûs* 34 (2000): 517-549.

Rorty, Richard. *Philosophy and the Mirror of Nature.* Princeton: Princeton University Press, 1979.

Russell, Bertrand. *The Problems Of Philosophy*, 1912.

Scanlon, T.M. *What We Owe To Each Other.* Cambridge: Harvard University Press, Belknap Press, 1998.

Schiffer, Stephen. "Contextualist Solutions to Scepticism." *Proceedings of the Aristotelian Society* 96 (1996): 317-333.

Searle, John. *Rationality in Action.* Cambridge: MIT Press, 2001.

Sosa, Ernest. "Epistemology Naturalized." In *Knowledge In Perspective: Selected Essays in Epistemology.* Cambridge: Cambridge University Press, 1991.

_____. "Knowledge and Intellectual Virtue." In *Knowledge in Perspective.*

_____. "How to Resolve the Pyrrhonian Problematic: a Lesson from Descartes." *Philosophical Studies* 85 (1997): 229-49.

_____. "Skepticism and Contextualism." *Philosophical Issues* 10 (2000): 1-18.

Stich, Steven P. "Reflective Equilibrium and Analytic Epistemology." In *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*, 75-100. Cambridge: MIT Press, 1993.

Stroud, Barry. *The Significance of Philosophical Scepticism*. Oxford: Clarendon Press, 1984.

_____. "Understanding Human Knowledge in General." In *Knowledge and Skepticism*, ed. M. Clay, K. Lehrer, 31-50. Boulder: Westview Press, 1989. Reprinted in *Knowledge: Readings in Contemporary Epistemology*, ed. Sven Bernecker and Fred Dretske, 307-323. Oxford: Oxford University Press, 2000. Page references are to this reprint.

Unger, Peter. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press, 1975.

Velleman, David. *Practical Reflection*. Princeton, NJ: Princeton University Press, 1989.

_____. *The Possibility of Practical Reasoning*. Oxford, New York: Oxford University Press, 2000.

Wallace, R. Jay. *Responsibility and the Moral Sentiments*. Cambridge: Harvard University Press, 1996.

_____. "Three Conceptions of Rational Agency." *Ethical Theory and Moral Practice* 2 (1999): 217-242.

Watson, Gary. "Free Agency." In *Free Will*, 96-110. Oxford: Oxford University Press, 1982.

Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich. "Normativity and Epistemic Intuitions." Forthcoming in *Philosophical Topics*. Available online at:

http://ruccs.rutgers.edu/ArchiveFolder/Research%20Group/Publications/NEI/NEIPT.html. Accessed February, 2002.

Williams, Bernard. "Deciding to Believe." In *Problems of the Self: Philosophical Papers 1956—1972*, 136-151. Cambridge: Cambridge University Press, 1973.

_____. "Moral Luck." *Proceedings of the Aristotelian Society* Supplementary Volume 50 (1976): 115-135.

Williams, Michael. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism.* Princeton: Princeton University Press, 1996.

_____. "Contextualism, Externalism and Epistemic Standards." *Philosophical Studies* 103 (2001): 1-23.

Wolf, Susan. *Freedom Within Reason*. New York: Oxford University Press, 1990.

Curriculum Vita

Jill Claudia Rusin was born October 24, 1968 in Chicago, Illinois. She received a BA, *cum laude*, in philosophy and science from the University of Pennsylvania in May, 1990, and an MA in philosophy from the Johns Hopkins University in May 1996. She has been a teaching assistant at Johns Hopkins in logic, philosophy of mind, and for general introductory courses. She received a Deans Teaching Fellowship to teach an epistemology course of her own design, "Scepticism, Philosophy, and 'Common Life'." Rusin has taught feminist philosophy at Hopkins and a graduate course in scientific reasoning at University of Maryland, Baltimore County. She has presented papers at Columbia University, the University of Rochester, Johns Hopkins University, and the Dominican College (Ottawa, Ontario). She also served as student representative on a departmental hiring committee, and as the president of the Hammond Society, Hopkins' philosophy graduate organization.