

CONTEXTUALISM AND SKEPTICISM: THE DEFEAT OF THE BOLD SKEPTIC

Contents

1. Contextualism and the Project of Solving the AI Puzzle.....	1
2. A General Puzzle about Skeptical Inclinations Waxing and Waning vs. the AI Puzzle I'm Trying to Solve.....	4
3. Unspoken AI	8
4. Actual, Spoken Disputes over AI: I'm Not that Nice Contextualist!.....	9
5. Are the Skeptic and Her Moorean Opponent Both Making True Claims as They Argue?: Some Disputes Are Genuine!.....	11
6. Are the Claims that Our Disputants Are 'Trying' to Make Both True?: Some Disputes Are <i>Deeply</i> Genuine!	14
7. Are the Claims that Our Disputants Are 'Trying' to Make Both Deeply Important? 16	
8. Bold vs. Timid Skeptics	17
9. Kornblith's Attack: Full-Blooded vs. High Standards Skeptics	18
10. The Philosophical Interest of the High Standards Skeptic	21
11. Do I Respond Only to the High Standards Skeptic?: The High Standards Skeptic, the Deplorable Position Skeptic, and the Bold Skeptic	26
12. Do I Respond Only to the High Standards Skeptic?: The Simultaneous Defeat of the Bold Skeptics of Both Kinds.....	29
13. My Supposedly Thoroughly Externalist Response to the Full-Blooded Skeptic 33	

14.	Indexed AI	36
15.	Irrelevant to Traditional Epistemological Reflection on Skepticism?	39
16.	Is My Contextualist Position Inexpressible?.....	43
17.	The Factivity Problem	47

CONTEXTUALISM AND SKEPTICISM: THE DEFEAT OF THE BOLD SKEPTIC

1. Contextualism and the Project of Solving the AI Puzzle

Does the contextualist seek to dissolve disputes over skepticism in such a way that all parties to them come out being right? And does she use a ‘perfectly general strategy’ for doing so? Is she ignoring the traditional epistemological topic of whether we really know things, instead addressing how the word ‘know’ is or should be used? Is she engaged in philosophy of language instead of epistemology? Is she addressing the more important types of skeptic? Are key aspects of her position inexpressible, by her own lights? Is she subject to a ‘factivity problem’?

These and other questions are answered in this chapter, as I explain and defend the contextualist nature of my solution to the puzzle that the skeptical argument (AI) presents us with. A clear view of just what the contextualist—or at least *this* contextualist—is trying to achieve, and how I am trying to achieve it, is crucial to most of the responses, tightly tying together the ‘explain’ and the ‘defend’ aspects of my efforts. I will be speaking for myself here, and there may be significant differences in approach among different contextualist responders to skepticism, but I do think that on issues pertaining to the contextualist nature of our responses to skepticism, I am for the most part at least roughly aligned with at least my fellow ‘contextualist Mooreans’ who were most influential on me: Gail Stine, Stewart Cohen, and David Lewis. While I will be taking positions on some specific questions that my fellow contextualists may not have answered, and while I can’t and don’t presume that they would give the same answers, I do think that the moves I make for the most part fit well with their general frameworks, and so are answers they

could well give.¹ While I'll be addressing concerns as they apply to my own contextualist resolution of AI, contextualist responses can be made to other skeptical arguments as well, where some of the same issues will arise, and where aspects of my treatment of AI can be applied. Moves analogous to some of those I make here are of potentially quite broad application.

I will be devoting my attention to the most important objections to contextualist solutions to skepticism (and to me in particular) that have been leveled in published criticisms,² but will also be presenting my defense with an eye to answering misconceptions of the contextualist approach to skepticism one is likely to encounter on one's philosophical travels. Of course, these types of resistance overlap as popular lines of criticism find their way into print, or as published criticisms find their way into the general philosophical consciousness.

One of the most extreme, and most common, worries about contextualist solutions is one perhaps most effectively expressed by Ernest Sosa: that they are simply irrelevant to traditional epistemological reflection on skepticism. Because this worry concerns how my approach relates to other philosophical approaches, it is best to hold off on responding to it in full until later in this chapter, when my approach has been more fully developed. However, it is worth noting here at the outset that those who have worries like Sosa's often seem to view contextualist responses to skepticism as engaged in a very different inquiry from the one that epistemologists are usually or traditionally involved in when they address skepticism. Contextualist and 'traditional' approaches are seen as being on opposite sides of some impressive-sounding divide or other. This can be tied up with issues of classification, of which sub-area of philosophy the projects are said to be part of: Contextualists are doing philosophy of language; traditionalists, epistemology. Or, relatedly, it can be seen as some kind of difference in levels of semantic ascent:

¹ The exception I have particularly in mind in connection with that 'for the most part': In section 13, I'm dealing with an issue that may pertain specifically to me in virtue of some details of the account of knowledge I have used in my work, and so is not relevant to other contextualist accounts.

² With this exception: One of the most important objections to contextualist responses to skepticism, and mine in particular, is that leveled in (Schiffer 1996), to the effect that our responses involve us in implausibly attributing 'semantic blindness' to speakers. However, since Schiffer's objection is so tightly bound up with important objections to contextualism itself, it was answered, in a way I am still quite happy to stand by, in vol. 1, at (DeRose 2009: 174-9).

Traditionalists are investigating questions concerning whether we know; contextualists, questions concerning the proper use of 'know(s)' (e.g., Sosa 2000).

So, perhaps the first thing to be said in characterizing the contextualist approach is that, as I view things, and so far as I can see, that's all wrong. As a general rule, when such a divide is posited, and the contextualist is said to be doing 'not *this*, but only *that*', what is really happening is that I *am* doing *that* in a way that involves *this*. So as I see things, I am addressing an old problem—the traditional, philosophical problem of skepticism, at least as it is raised by one venerable form of skeptical argument. And though I am indeed bringing in some (fairly low-tech) philosophy of language, and some talk about the truth-conditions of sentences containing 'know(s)', I am not doing any of that *instead of* addressing the question of whether we know things in light of AI-like skeptical arguments. Rather, I am bringing in the relevant philosophy of language *as a way of* addressing the same (epistemological) issue that has been addressed by other approaches.

This is all best viewed in light of the methodology discussed a couple of chapters ago. As we saw in Chapter 2 (expanding on the methodological remarks in SSP), I employ a Moorean methodological approach to such skeptical arguments, seeing the problem as one of a set of claims, each of which can seem very plausible, but which seem mutually inconsistent with one another. So far, of course, that does not constitute changing problems, but just taking a sensible methodological approach to an old problem. Going beyond Moore himself (so far as I can see), I take the most promising way forward in adjudicating the conflict to be in the construction of a solution to the problem which provides a good ('damage-controlling') explanation for why we go wrong in our thinking at the points where the solution charges us with error. But this is still just sensible Moore-inspired methodology, even if not the methodology of Moore himself, and does not constitute any big change in the very nature of the problem being addressed.

It's here that I bring in contextualism: It, together with a partial account of what knowledge is (a 'double-safety' account, as I've come to call it since the writing of SSP, and that will be explained more fully in Chapter 7) provide the materials for the explanation I propose. But these are just the tools I use to do the job of solving the problem; other would-be solvers of this very same problem are welcome to use other tools, and to not use contextualism, nor any philosophy of language at all, in their attempts. If I judge such an alternative solution to be a (relative and/or absolute) failure, it won't be because I insist at the outset that any good solution

must use tools like those that I use, but because I judge that, as it turns out, the alternative just does not succeed in solving the (same, old) problem—the one that I too am claiming to solve.

As noted in the first two chapters, my contextualist solution does differ importantly from the typical ‘straightforward’ solutions in that, by offering different readings of the three claims constitutive of our puzzle (AI’s two premises and the negation of its conclusion), it distributes the intuitive costs among more than one of the three claims constitutive of our puzzle before it seeks to explain the problems away: Each the first premise of AI and the negation of AI’s conclusion are true on one of the salient readings of it, false on the other. (The second premise is true on any standard for knowledge, so long as that standard is held constant between the two occurrences of ‘know’.) But I don’t insist that is how things must be done: ‘straightforward’ (non-contextualist) solutions are welcome to choose just one loser here, and explain away why we can find it so plausible, though it is, according to the solution, wrong. Indeed, a main part of SSP (sects. 15-16) is arguing that attempts to provide just such damage-controlling explanations on behalf of bold skepticism (which chooses the negation of AI’s conclusion as the straightforward, lone loser), despite some initial promise, *turn out* not to succeed.

2. A General Puzzle about Skeptical Inclinations Waxing and Waning vs. the AI Puzzle I’m Trying to Solve

In this telling statement, Timothy Williamson touches on an area of confusion about the nature of contextualist responses to skepticism that it will be important for us to clear up:

Contextualism supplies a perfectly general strategy for resolving any apparent disagreement whatsoever. Since some disagreements are genuine, we should not always follow that strategy. The conflict of intuitions does not always disappear on further reflection. At least some intuitions are mistaken. Moreover, they are explicably, not blankly, mistaken. (2005: 694-5)

I hope the previous section, and, even more so, the previous two chapters, make clear (or renders even clearer what was already clear in SSP) that I see myself very much in the game of addressing the conflict of intuitions that at least one skeptical

argument presses on us, and that I recognize attempts by non-contextualists to explain how some of the intuitions involved are simply mistaken as a perfectly fine way in principle to address such conflicts.

The important area of potential confusion I have in mind is one that surrounds the assumption, made here by Williamson, that contextualist solutions don't make disputes over skepticism out to be genuine disagreements. I'm not sure what exactly Williamson means by 'genuine' here, but often when it's said that contextualist solutions render disputes over skepticism non-genuine, what's being alleged is that the contextualist construes the judgments being expressed by the parties to such a dispute to be compatible with one another (and, in fact, to both be true). At any rate, just when and in which ways contextualism renders disagreement over skepticism genuine turns out to be a somewhat complicated matter that I will be addressing in the sections that immediately follow this one. My focus will be on disagreements that occur between AI-wielding philosophical skeptics and Moore-like resisters of skepticism, because such disagreements are the ones most germane to my project, and ones about which some of the most serious misunderstandings of the contextualist approach arise.

However, before addressing such disagreements, I should note that the material around the above-quoted passage seems to indicate that Williamson also has his eyes on a kind of disagreement (or at least a *difference*, if one doesn't think it proper to use 'disagreement' here) in intuitions that might be had by a single individual at different times over a change in settings: In an epistemology seminar in which skeptical worries prevail, one might have skeptical feelings, inclinations, or intuitions to the effect that she does not know that O, while the same person may quite firmly judge or intuit that she does know that O later, in a more ordinary setting. This general phenomenon, the skeptical side of which can be driven by a variety of different skeptical arguments, is frequently noted, sometimes citing Hume's famous observations about how skeptical doubts felt in philosophical contexts dissipate in other settings.³ And some may think contextualism is supposed

³ Hume, *A Treatise of Human Nature*, Book I, Part IV, section 7 (Hume 1978: 269): 'Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium, either by relaxing this bent of mind, or by some avocation, and lively impression of my senses, which obliterate all these chimeras. I dine, I play a game of back-gammon, I converse, and am merry with my friends; and when after three or four hour's amusement, I wou'd return to these speculations, they appear so cold, and strain' d and ridiculous, that I cannot find in my

to be primarily supported by its ability to explain how such shifts in intuitions and leanings can occur, or that contextualists themselves seek to support their views by such considerations. And we contextualists have no doubt thought and written things to encourage such an understanding of what we're up to (as, for instance, in section 2 of SSP). However, even getting a decent grip on the phenomena to be explained here (just when such shifts occur and what these shifts are like) turns out to be quite tricky, I think, as this all seems to work out very differently for different people.⁴

heart to enter into them any farther.' Crispin Wright prominently alludes to this passage at (Wright 2005: 240).

⁴ Williamson points out that a non-contextualist explanation of such shifts in an individual over time is possible:

Whether one knows (or one's belief is epistemically justified) also depends on the balance between many complex, unquantifiable, subtly interacting considerations. In many realistic cases, some weigh one way, others the other way. The concept of knowledge or of epistemic justification provides no algorithm for weighing all these factors against each other or integrating them into a final verdict. No wonder that the skilled sceptic can present the considerations that favour a negative verdict so vividly that they intuitively appear to outweigh the considerations on the other side. It does not follow that the sceptic is right, even in the context of the epistemology seminar; the case may not even be borderline. Nor does it follow that the sceptic is wrong, even in the everyday context. (2005: 695)

In fact, Williamson goes on to argue that an important feature of such shifts favor a non-contextualist account:

As before, the intuitions that predominate in one context spill over to judgements about the truth-values of sentences as uttered in the other context. In the everyday context, it is intuitive that someone in the epistemology class who says 'Mary does not know that the train leaves at noon' is overestimating her epistemic difficulties. In the epistemology class, it is intuitive that someone in the everyday context who says 'Mary knows that the train leaves at noon' is underestimating her epistemic difficulties. Although such data are not decisive against contextualism, they tend to support the non-contextualist explanation. (2005: 695)

For the record, I think Williamson's statement about what the 'data' are here is overly simple: In either setting, if we have the ground-level intuitions Williamson assigns to us, when we then turn our thoughts to the truth values of our claims made in the other setting, which is in relevant ways very different from the setting from which we are making the judgment about the claim, I believe that things get very murky to us very quickly—and in my experience different people say very different things. Further, I think that the murkier, actual data about how in one setting we regard the judgments we make in very different settings would be more favorable to contextualism playing an important role in explaining what's going on than are the much cleaner 'data' that Williamson seems to suppose—though it is likely that other forces, having little to do with contextualism, are importantly at work here as well.

But there are important grounds for a contextualist solution to AI-inspired skepticism that do not rest on its ability to account for any of the fine (and apparently quite variable) details of how skeptical doubts, however they might arise, wax and wane in their force over time as one moves from one setting to another. We should fix our eyes on a few facts about our reactions to the constitutive parts of the AI puzzle, taken individually: At least when the claims are considered individually, we do tend to find the claim that we know that O very plausible, but we also can (see the discussion at esp. sections 7-10 of Chapter 2) find each of the skeptic's premises very plausible. That's very puzzling, to say the least, given the relation among these claims. *That's* our AI puzzle, which I have been out to solve. Even if, quite contrary to fact (insofar as I've been able to determine the relevant facts), say, almost everyone were very strongly inclined to accept AI's conclusion when presented with that skeptical argument, there would still be the puzzle of why we are so strongly inclined to judge that we do know that O in other settings, even though we are also inclined to be attracted to each of AI's premises. And what I claim to explain is how AI's premises can have the intuitive appeal that they have, even though its conclusion can (at the very least, in most settings) seem so incredible.

Of course, though I have been focused on one particularly strong basis for skepticism, contextualist solutions in important ways analogous to my solution to AI may also be appropriate for dealing with other, quite different, skeptical arguments. But if so, and if these solutions really are in the significant way analogous to mine, the important question will be whether the proposed contextualist solutions accomplish what I claim to accomplish with respect to AI: Can the contextualist semantics of 'know(s)' being proposed (supposing the skeptical argument in question concerns knowledge), together with something like a (no doubt partial) account of what knowledge is, well explain, and explain better than rival accounts, how the premises of the particular skeptical argument being addressed can be as plausible as they are, while its conclusion is as implausible as it is? *We certainly* should not be trafficking in anything that could with any fairness be called a one-size-fits-all 'perfectly general strategy for resolving any apparent disagreement whatsoever.'⁵

⁵ There is something of a shallow recipe by which one might use contextualism (or also, I suppose, for that matter, bare, unsupported claims of two-way ambiguity) to, at least by some low standard, 'resolve' any

3. Unspoken AI

It can seem strange, or even wrongheaded, for the contextualist to appeal to rules of conversation for the setting of epistemic standards in trying to solve the AI puzzle, for, though a real, live skeptic could present AI to another person she is actually talking to, the conundrum AI confronts us with seems just as puzzling when one instead wrestles with it in solitude, without anyone saying anything at all. How could a rule *of conversation* then play such a pivotal role in its solution? To revisit and expand upon a point made in section 2 of SSP (the third ‘important point’ made at the end of that section), the key to the answer to that pointed question is that our judgment as to whether we know something, even when that judgment is issued silently and while in solitude, seems to be influenced by our sense of whether it would be right or wrong to claim to ‘know’ it. Central to my solution is my contention that, due to the operation of the Rule of Sensitivity, claims to know that not-H and admissions that one does not know that not-H tend to invoke high epistemic standards that one’s belief that not-H does not satisfy, which produces a tendency for the claim that one does ‘know’ to go false and the admission that one does not ‘know’ to be true. That the claim has such a tendency to backfire, and the admission to work out so well, can make it seem to us, even when we’re considering the matter silently and in solitude, that we really don’t know that not-H.

Some such move also seems needed to make sense of various broadly Gricean maneuvers, including some we should be keen to endorse. So, for example, insofar

apparent disagreement whatsoever. Here is Stanley’s way of putting it (though he is describing the contextualist’s handling of Bank-like cases, rather than its handling of skepticism): ‘Suppose we have two claims, *S* and *not-S*. It always open to someone to resolve the apparent conflict by maintaining that the two occurrences of ‘*S*’ in the claims express different propositions, relative to their differing contexts of use’ (2005a: 122). But the actual proposed contextualist solutions to skepticism are not just instances of that recipe. Or at least no more so than the actually proposed invariantist approaches to the problem are just instances of the very general invariantist ‘recipe’: ‘Insist that the two occurrences of “*S*” in the claims always express the same proposition in all contexts, so that the two claims are simply incompatible, and then deny one of them.’ That ‘recipe’ seems to be one that, for all we can see in advance, could be followed very effectively or extremely ineffectively (depending crucially on how good a case is made for denying the claim that is rejected). Well, of course, the same goes for contextualist solutions. Perhaps Williamson would agree with all that. Still, it needs to be said.

as you were tempted to think that, say, a bright red mailbox that you are getting a very good look at from only two feet away and in bright sunlight does not look red to you, on such grounds as that it is not doubtful enough to you that the object is red for it to look red to you ('It doesn't *look* red to me; it very obviously *is* red!'),⁶ you would likely be so tempted even when thinking through the issue silently and in solitude while staring directly at such a mailbox, or imagining so staring, and not only when actually engaged in making the relevant assertion to some listener. That point doesn't nullify Grice's attempt to explain away the temptation (Grice 1961). ('How can a rule of *conversation* possibly explain away why 'It looks red to me' can seem false? The appearance of falsehood is there just the same even when I just think the sentence to myself, with nobody *saying* anything!') Grice's maneuvers, and other maneuvers of roughly similar intent, should be understood as directly explaining what's wrong with the spoken claim (in the relevant circumstances), but then also indirectly explaining why we may be tempted by a false view of the meaning of 'looks' (or whatever) even when thinking to ourselves: Our sense that it would be wrong to make the claim in the relevant circumstances can tempt us toward (wrongly) judging that the claim would there be false (where one has no doubt the object referred to is red), even when we make the judgment silently and to ourselves, and then to a wrong account of the meaning of the claim.

4. Actual, Spoken Disputes over AI: I'm Not that Nice Contextualist!

Though I don't think the matter of which side of various unenlightened (as we're about to call them) debates over skepticism may have 'won' is very important (see DeRose 2009: 138-9), some folks may have actually taken part in such disputes, and may be interested in whether it was they or their opponent whose contentions were correct. But there is another, potentially more important, reason, beyond this matter of personal pride, that some have for being interested in what the contextualist has to say about such disputes: They may think that the contextualist's ruling about them is implausible, and so counts against the contextualist approach to skepticism.

⁶ I realize that in these post-Gricean times, many of us are not tempted very far by such thoughts.

So, let's consider such a dispute.⁷ Imagine a character we can call an 'undifferentiated' AI skeptic. Sadly, the possibility of contextualism has not occurred to her, even as a position someone else might accept, so she just takes it for granted that we either simply know that O or we simply don't know that O;⁸ and on the basis of AI she holds that it's the latter of those two that is the case.

Now we arrange for our skeptic to enter into a spoken dispute. So, suppose this undifferentiated skeptic meets resistance as she presents AI. The possibility of contextualism has also not occurred to the undifferentiated Moore-like resister she encounters, who also presumes that we either simply know that O or simply don't know that O, but who, for roughly Moorean (methodological) reasons, insists in the face of our skeptic's argument that it's the former of those two that gets things right. 'How absurd it is to suggest that I do not know it, but only believe it, and that perhaps it is not the case!', he exclaims.⁹ Back and forth these combatants argue in a heated, but as-yet 'unenlightened', disagreement—i.e., a disagreement sadly unenlightened by the possibility of contextualism.

Enter the nice contextualist. This character is apparently how contextualists are often imagined to be. He is a very affirming, perhaps even nurturing, soul. Like a parent who seeks to intervene only in neutral, balanced ways in squabbles among his children, whenever this nice contextualist encounters a philosophical dispute, or at least one meeting certain conditions (perhaps one that seems difficult to a certain degree to adjudicate), he seeks to 'dissolve' it as being only an apparent disagreement, in such a way that the parties to the dispute are all construed as making true claims as they argue, and/or, relatedly, he seeks 'even-handed'

⁷ In addressing what to say about such a dispute over skepticism, I take myself to be addressing an issue that looms large in popular conceptions about what contextualists are up to. But criticisms of the contextualist approach to skepticism that start with thoughts about what the contextualist says about such disputes are in the published literature. As I point out in Volume 1 (DeRose 2009: 130-131), a particularly good example of this is (Richard 2004: 215-16).

⁸ Or she may also add the possibility that it's simply indeterminate whether we know: that claims to the effect that we know the various Os are neither true nor false. Recall that on our semi-technical use of 'simply', 'We simply know or we simply don't know' isn't used to assert bivalence about these matters, but to claim that the matter of what the correct answer is doesn't vary with contextual shifts in the meaning of 'know(s)'.

⁹ This exclamation is based on what Moore said in his famous 'Proof of an External World' in defense of the claim that he knew that 'Here is one hand and here is another': 'How absurd it would be to suggest that I did not know it, but only believed it, and that perhaps it was not the case!' (Moore 1959a: 146).

resolutions where the positions of all the parties to the dispute are respected and validated to (at least roughly) the same degree.¹⁰ We might further suppose that it was largely in order to facilitate such a nice resolution of such often nasty disputes over skepticism that our nice character was led to accept contextualism in the first place.

But that ain't me, as we're about to see. Nor do other actual contextualists about knowledge and skepticism tend to be like that, as we'll also note.

5. Are the Skeptic and Her Moorean Opponent Both Making True Claims as They Argue?: Some Disputes Are Genuine!

Contextualism does open up the *possibility* of ruling that both of our unenlightened disputants are making true claims as they argue. But what do actual contextualists say about such a dispute? Probably the first thing to say is that we generally haven't cared all that much about the question—and we certainly aren't following any general imperative to rule that both parties to such a dispute are making true

¹⁰ I am inspired in my use of 'even-handed' here to characterize how many philosophers seem to think of the motivations behind contextualist solutions to skepticism by (Wright 2005: 240-242), but Wright himself uses the term to describe, not how we contextualists seek to treat the two sides of philosophical disputes over skepticism (philosophical skeptic vs. Moore-like resister of skepticism), but how we seek to treat skeptical denials of knowledge vs. commonsense affirmations of knowledge (philosophical skeptic vs. ordinary claims made in non-philosophical settings). In so viewing us, Wright is certainly getting closer to our (or at least my) actual motivations than is the common conception of contextualists I am using his term to characterize. However, Wright seems to me to still be taking us to be more concerned with the fate of the skeptic than I think we really are. As I see it, when we contextualists set out to address the puzzle that the skeptic's argument confronts us with, we were seeking to adequately account for the power of the skeptic's argument, but we were more open than Wright takes us to be to solutions on which the skeptic turned out to be just mistaken. But, of course, judging motivations—including, and perhaps especially, one's own—is delicate business! Still, I especially find the portion of the following passage from Wright to which I'm adding emphasis to be inaccurate, at least as a description of my aims: 'A perfectly even-handed treatment of scepticism and common sense will allow their respective conclusions to stand relative to context *and will charge neither with any cognitive shortcoming unless the charge is matched by a corresponding charge against the other*. So the claims of common sense to knowledge will stand, albeit defeasibly, no doubt, modulo certain relevant contextual factors. And the claims of scepticism to unavoidable ignorance will also stand, modulo certain relevant contextual factors' (Wright 2005: 241).

claims. I reject alternative solutions because they turn out not to work, not out of hand because they fail to deliver the ‘nice’ verdict that both parties to the dispute are speaking truthfully as they argue.¹¹

Indeed, I *could* not reject those bold skeptical solutions for that reason, because, like other actual contextualists, I myself don’t opt for such a ‘nice’ solution on which both parties come out making true claims! However, I did not explain that aspect of my view in SSP itself, but there provisionally assumed a skeptic-friendly version of contextualism for ease of exposition (see the second ‘important point’ explained in section 2 of SSP), waiting until ‘Single Scoreboard Semantics’ (‘SSS’)¹² to say how I was inclined to address those issues. So a quick review of the situation is called for here.

Whether a contextualist will rule that our arguing parties are both making true claims depends on what not-so-obvious call he makes on the general issue of how to handle a kind of dispute in which: a context-sensitive term is crucially involved; the disagreeing parties are making claims that at least seem on the surface to be in conflict with one another; whether they think about what they’re doing in these terms or not, each arguing party is executing conversational maneuvers that tend to push the content of the key context-sensitive term toward having a value that would make their own claims true;¹³ and yet they each take themselves to be

¹¹ I do hold it against bold skeptical solutions that they involve us in ‘systematic falsehood’ in our use of ‘know(s)’ (see esp. section 16 of SSP), but this is general systematic falsehood in speakers’ use of the term, importantly including uses I am much more hesitant to rule are false, as compared with some falsehoods issuing from the mouths of some Moore-like disputants in the odd philosophical dispute over skepticism (see DeRose 2009: 56-59). And I put this forward as the kind of problem that any solution to a genuine conflict of intuitions must face and try to overcome. I reject the proposed bold skeptical solutions because their attempts to overcome their problem turn out not to work as well as their contextualist rival in making sense of the various intuitive claims we are, or at least should be, trying to adjudicate among in navigating this puzzle. Importantly to our current concern, I definitely do not reject them out of hand because they fail to deliver the ‘nice’ result that both parties to our undifferentiated dispute are making true claims.

¹² SSS was first a paper (DeRose 2004a), then Chapter 4 of (DeRose 2009). Though I’ll be giving page references to the latter, and though the latter does contain important material not included in the earlier version, everything I use and refer to here was also in the earlier version.

¹³ One important type of mechanism for determining the content of context-sensitive terms seems to be ‘rules of accommodation’ on which the content takes on values so as to make what is said true. If so, then it is to be expected that the claims of each party to such a dispute will tend to push the conversational score toward values that makes that party’s claims true.

contradicting one another, and may even give explicit indications of their intent to be contradicting one another ('No, you're wrong. Sally *is* tall'). A contextualist who seeks to make both parties to such dispute come out speaking truths will opt for something like what I call a 'multiple scoreboards' view on which the content of each speaker's spoken claims matches that speaker's 'personally indicated content'—the content that the conversational maneuvers that speaker is making would tend to put into place (DeRose 2009: 134-5). And many assume that is precisely what the contextualist should say here, about our imagined dispute over skepticism, or that actual contextualists do take that option, or even that it is a desire to attain the nice result that both parties are making true claims that motivates contextualism in the first place.

But that must not be what's really motivating us, because none of Gail Stine, Stewart Cohen, David Lewis, or myself is inclined toward such 'multiple scoreboard' contextualist views, but instead opt for views on which such a dispute is quite real and it's impossible for both unenlightened disputants to be speaking truthfully.¹⁴ On such 'single scoreboard' views, the 'conversational score' for the term in question is affected by the conversational moves made by both parties to the argument, but the result of all that maneuvering is a single 'score' that governs the content of both parties' use of the term, so that the one is denying just what the other is affirming (2009: 135-6). In my own case (and I suspect in the case of Stine, Cohen, and Lewis), a main source of attraction to such a 'single scoreboard' approach is that I very much share the (often strong) sense our disputants are contradicting one another, and the approach makes good sense of that feeling. That basic approach, though, leaves open the questions of what that single score is and so which, if either, of the parties is making true claims. These are matters that we contextualists haven't found very urgent, but, as it happens, on the particular view that I lean toward, and on both of the views that seem to be Cohen's finalists, *neither* the skeptic nor her opponent is speaking truthfully as they argue. Rather, both parties to the dispute are making claims that are neither true nor false (2009: 58, 144-8). Lewis tends toward a more skeptic-friendly version of contextualism, on which the skeptic is speaking truthfully in such an argument, but Lewis seems no more anxious for a 'make

¹⁴ Stine's relevant pioneering writings on these matters was quite early, and, due to her tragic early death, she didn't face these issues very explicitly, so my classification of her must be somewhat tentative, but I am inclined to read her (in Stine 1976) as tending toward (in the terminology of SSS) a single-scoreboard view of the 'reasonableness' variety.

everybody happy' solution than Cohen or I am, for on his view, the skeptic's opponent is making false claims.¹⁵ And though I 'provisionally assumed' a certain position on the matter in SSP for ease of exposition and not because I thought it was correct, even that view was one on which only the skeptic is making true claims in the type of dispute in question. So, though contextualism opens up the possibility for us to do so, none of us seems to opt for a view that delivers the 'nice' result that both parties to such a dispute are speaking truthfully as they argue.

Of course, 'punishing' both of the disputants, as it were, by ruling that neither of them is making true claims, as Cohen and I are inclined to do, introduces an element of 'even-handedness' to our treatments, in the manner of a certain kind of evenly harsh parent. But it does clearly show Cohen's and my willingness to depart, I suppose in just about as radical a way as possible, from the kind of 'everybody's happy' (in the 'everybody's making true claims' sense of that) solution the achievement of which motivates the contextualist of the imaginations of many.

6. Are the Claims that Our Disputants Are '*Trying*' to Make Both True?: Some Disputes Are *Deeply* Genuine!

For the contextualist (one who follows the 'basic contextualist strategy' of section 2 of SSP), there are truths in the vicinity of what both of the disputants in our unenlightened argument over skepticism are saying, and these truths can be thought of as what the disputants are 'trying' to express, as it were: We do know that O by ordinary, moderate standards for knowledge that the Moorean's resistant maneuvers would tend to keep in place (or put in place, depending on details of the story), and we do not know that O by the standards the skeptic's maneuvers have at least some tendency to put in place. Given certain details of how the contextualist story is told, one or both of our disputants may be blocking the other from making the true claims they are in this way 'trying' to make—i.e., the true claims they would be making if their conversational-score-affecting moves were accepted or went unopposed (and on invariantist theories, where the score doesn't change, we say

¹⁵ See (DeRose 2009: 136-8) for discussion.

that one is ‘trying’ to say just what one does say).¹⁶ This is the feature of my contextualist solution that those who take me to be a ‘nice’ contextualist are likely picking up on, and they might respond to what I say in the previous section as follows:

OK, so you evidently don’t seek to make both parties of our dispute over skepticism out to be making true claims. But, still, you do insist on finding two true propositions, one of which can be expressed by the sentence ‘I know that I have hands’, and the other by ‘I don’t know that I have hands’, and though you don’t rule that our disputants are each succeeding in asserting one of these truths, you do make them out to be each ‘trying’ to express one of them, as it were, in the way explained above. So you really are, in a quite relevant way, seeking some kind of ‘no-fault’ dissolution to the disagreement, in which there is a truth that is at least in an important way standing behind each of the disputant’s positions. And it is a mistake to insist on such a solution to just any philosophical dispute. Some disputes are, as we may say, ‘*deeply* genuine’: the disputants cannot both be right, neither on the level of what they are actually claiming, nor on the level of what they are ‘trying’ to claim, as we have been using that phrase. And perhaps the dispute over skepticism is one of these deeply genuine ones.

My solution does indeed differ from invariantist solutions in that it does not make the disputes over skepticism in question out to be deeply genuine in this way. And I suppose that some easy-enough-to-imagine philosopher could start their inquiry into AI-inspired skepticism with some kind of sense that the dispute here is not (in the peculiar sense we are now using) deeply genuine, and might be strongly motivated toward contextualism and a contextualist solution to skepticism precisely because they vindicate that sense. In an extreme case, such an inquirer could be well described as ‘insisting’ on a solution that delivers that result—and that could be their inclination toward philosophical puzzles and conflicts generally, or perhaps toward those that meet some general condition, like being sufficiently difficult to resolve.

But that need not be what motivates contextualism, and that’s not what motivates me, so far as I can tell. I reject invariantist solutions to AI, including the bold skeptic’s solution, because they *turn out* not to work as well in resolving the puzzle as does the contextualist solution I do accept, not because I ‘insist’ from the

¹⁶ Here, I employ, while being more explicit about my intended meaning, a use of ‘trying’ that I used back in (DeRose 1990: 275).

outset that any successful solution to this problem—nor *certainly* toward just *any* sufficiently knotty philosophical problem—will have to make the dispute in question out to be not deeply genuine.¹⁷

7. Are the Claims that Our Disputants Are ‘Trying’ to Make Both Deeply Important?

Some understand contextualist solutions as seeking to provide a response to the AI dispute that is ‘even-handed’ in that it acknowledges the deep importance of both of the truths that our arguing parties are ‘trying’ to express, like a lecturing parent: ‘Now, Abby, you are trying to make a very important point here. But what Ben is trying to say is equally important. You two should just stop all your fighting, allow each other to make your important points, and respect the importance of what the other is expressing.’ A nice contextualist like this, happening upon an argument between a skeptic and a Moore-like resister, occurring in dark ignorance of the possibility of contextualism, would, in bringing contextualist enlightenment to this fight, counsel each disputant to perhaps append some kind of clarifying phrase to their uses of ‘know’, stop presenting their claims as if they were at odds with what the other is saying, and achieve a blessed harmony in which each is making the true and deeply important claim that had been all along behind the value of the position they were taking.

But the basic contextualist strategy is neutral about the attitude one should take toward the truths in question. Though it would not be surprising to find someone attracted to contextualist solutions who finds the truth ‘behind’ the skeptic’s position to be very important (perhaps some deep, or even menacing, fact about the human condition), my sense is that we actual contextualists generally don’t think the skeptical truth in play here really is very significant. That is certainly true of me, as we will see in section 15, below. But the important point to make here is that contextualists need not be so motivated. The conviction that there’s bound to

¹⁷ However, I do hereby reiterate my admission from note 10 about the delicateness of judging one’s own motivations.

be something important to learn about knowledge from a puzzle as sharp as is the one presented by AI need not be driven by an opinion that there is some deeply important (and perhaps menacing) skeptical truth somehow standing behind the skeptic's position. The lessons to be learned may instead have to do with what knowledge is, and how knowledge claims work, in such a way as to *dodge* menacing forms of skepticism.

8. Bold vs. Timid Skeptics

So far, we've been keeping our AI skeptic in the dark, unenlightened by the thought of contextualism. Once the possibility of contextualism is raised, our AI skeptic faces a choice: Does she take her argument to show that we don't know even by the standards for knowledge that govern our ordinary, nonphilosophical epistemic talk, or only that we don't know by unusually elevated standards that her argument perhaps has some tendency to put into place? This is the choice between being a *bold* and being a *timid* AI skeptic, central to the divide-and-conquer strategy that I point to when I distinguish between these two skeptics in section 2 of SSP.

Our question to the AI skeptic does not presuppose contextualism. Our bold skeptic could actually be a contextualist who accepts varying standards for knowledge, but who thinks that AI's premises, and so its conclusion, are true by ordinary standards—perhaps (but not necessarily) because she thinks they're true by *all* allowable standards. But importantly, she also may well be an invariantist who thinks AI's premises are true by ordinary standards because they're true by the single set of standards that ever provide the truth-conditions for claims to 'know'—which ordinary standards she may think turn out to be much more demanding than a cursory look at our rather breezy knowledge-ascribing behavior might lead some to think. Or she could be undecided between the above: 'I'm not sure about contextualism, but in any case, my premises are true when evaluated at the standards that govern our ordinary talk'.

The timid skeptic accepts contextualism and takes her argument only to be showing that we don't know by the elevated standards for knowledge that the presentation of her argument has some tendency to put in place and not to be showing that we don't meet the standards that provide the truth-conditions for

ordinary uses of 'know(s)'. She may allow that we do know by ordinary standards for knowledge, or, more intriguingly, she may leave that question somehow open—perhaps suggesting something like that it must be left open because there's just no telling whether we know by ordinary standards.

These two skeptics face very different challenges, as I see things. The main challenge to the bold skeptic concerns whether and why (or *why in the world*, as the unsympathetic may put it) we should accept her bold skeptical claims, and pressing this challenge is my anti-skeptical focus in SSP. This is the skeptic I seek to *defeat*. Things are of course very different when it comes to the timid skeptic, for I *agree* with her defining claim that we don't know things by the high standards that the presentation of AI has at least some tendency to put into place. *Her* main challenge concerns the skeptical interest of her position: 'Whether such a timid skeptical stance is of any interest is a topic for another paper', I write in SSP, pointing to, but not really pressing, her main problem.¹⁸ As we will see in section 15, I in fact don't think the timid skeptic's true claims are very important, and so I don't take the 'even-handed' contextualist position that the undifferentiated skeptic (who has not yet considered the possibility of contextualism) is 'trying', as it were, to make a very important point that we should show due reverence to.

9. Kornblith's Attack: Full-Blooded vs. High Standards Skeptics

Hilary Kornblith's 'The Contextualist Evasion of Epistemology' (Kornblith 2000) provides a nice example of some of the kinds of misunderstandings that the contextualist approach to skepticism is subject to—and it has been quite

¹⁸ That 'other paper' still does not exist. It would/will be a reworking of Chapter 4 ('Wittgenstein's Suspicion and the Insignificance of Philosophical Scepticism') of my dissertation (DeRose 1990)—which, despite being perhaps my favorite chapter of my dissertation, is the one chapter that I still have not made into any publication. Some of the position I take—though not much of the thinking behind it—come out in section 15 of this chapter.

influential.¹⁹ It will hopefully function as a good target for pulling some doggedly persistent objections to contextualist approaches to skepticism out by the root. I will respond in detail (at spots, I suppose, exhausting detail) to Kornblith's attack, but mostly for the reason that doing so is a good way of clarifying my contextualist answer to skepticism in a way that should help in avoiding a wide range of misconstruals of it.

Kornblith minces no words. His paper is a follow-up to (Sosa 2000) where Sosa expresses his worry that contextualist solutions to skepticism are simply irrelevant to traditional epistemological reflection on skepticism. But, as we'll see (in section 15), Sosa, true to his judicious nature, does not claim that contextualist solutions are irrelevant; he just provides grounds for questioning their relevance and then wonders what the relevance might be. Kornblith, who takes Sosa to be 'extremely charitable in his description of the epistemological import of DeRose's contextualism,' enthusiastically takes up the 'less cautious' stance that 'DeRose's contextualism is irrelevant to epistemology' (Kornblith 2000: 25), and explains why he thinks that is so.

Central to Kornblith's case is a division he draws between two types of skeptics, which is like my distinction in SSP (explained in the section above) in that it supposedly features one skeptic who is both more ambitious and interesting than the other: his 'Full-Blooded' and his 'High Standards' skeptic. We will describe Kornblith's skeptics, as well as his eccentric Vermonter whom you're about to meet, in more detail soon. But note for now that the basic shape of Kornblith's complaint

¹⁹ Though this does not affect the main disagreements between us that I'm about to explore, I should note that Kornblith seems also to be among those who think of me as the 'nice' contextualist I've been distinguishing myself from in the previous sections of this chapter. Kornblith's paper opens with this:

Keith DeRose's contextualism offers both good news and bad news for the skeptic. The good news is that when the skeptic utters the words, "I don't know anything about the external world, and neither do you," the skeptic is making a true claim. The bad news is that when the anti-skeptic utters the words, "I do know a great deal about the external world and so does almost everyone else," the anti-skeptic is making a true claim as well. Put somewhat differently: the claim the skeptic makes is irrelevant to the claim the anti-skeptic makes. (Kornblith 2000: 24)

It seems that Kornblith is imagining his skeptic and his anti-skeptic as arguing with each other (based on the skeptic's use of the second-person pronoun and then the anti-skeptic's 'I do know,' instead of simply 'I know'), in which case, as we've seen, I actually do not rule that both speakers are asserting truths.

is that I answer only his 'High Standards Skeptic', the supposedly (completely) uninteresting of his two skeptics:

Now my worry about DeRose's contextualism can be put quite succinctly: DeRose responds only to the High Standards Skeptic, the skeptic who is an analog of my charmingly eccentric Vermonter and who acknowledges the importance and accuracy of substantive epistemological distinctions we wish to make. But this is a wholly trivial and uninteresting position. On the other hand, contextualism does nothing to address the Full-Blooded Skeptic, the skeptic who wishes to insist that all propositions about the external world are epistemologically on a par. But it is this latter skeptic who is making an historically important and philosophically interesting claim. If skepticism is a position we need to worry about, it is Full-Blooded Skepticism, not High Standards Skepticism, which should concern us. The skeptical problem DeRose's contextualism addresses is simply uninteresting from an epistemological point of view. (Kornblith 2000: 27)

While Kornblith's high standards skeptic is perhaps *relatively* uninteresting, she is also distinct from the 'bold skeptic' I explicitly take as my main target in SSP. As Kornblith does not mention my distinction between the 'bold' and the 'timid' skeptic, much less how his distinction relates to mine, I am unsure whether Kornblith simply misunderstands the nature of my announced target (thinking he is drawing the same distinction as me, but just using different labels), or whether he's aware of the difference between my 'bold' and his 'high standards' skeptic, but thinks that, despite my intentions, my response in fact misses its intended target and hits only the less interesting skeptic he describes. We will proceed in a way that should answer Kornblith's concerns in either case, and that will make use of his skeptics, as well as other skeptics that his discussion will prompt me to add to our cast of characters, to make clear where the importance and interest of contextualist responses to skepticism lies. We will start by looking at the considerable philosophical interest I at least think is to be found even in Kornblith's high standards skeptic; we will then see how my bold skeptic differs from the high standards skeptic in a way that adds still more interest to her position; and then we will see how my contextualist response really does answer the bold skeptic.

10. The Philosophical Interest of the High Standards Skeptic

Kornblith's 'Full-Blooded skeptic' is quite true to her name (in fact, *I* would be tempted to call her something like the 'wild-eyed' skeptic), holding that 'we are no more justified in believing that there is an external world than that there isn't' and that 'I have no degree of justification whatever for my claims about the external world. None' (Kornblith 2000: 26). By contrast, here is his supposedly uninteresting 'High Standards' skeptic:

If the skeptic agreed with us that there are differences in degree between the extent of justification which we have for various claims about the external world, but simply denied that we ever reach some very high standard required for knowledge, then skepticism would be a much less interesting claim. Indeed, imagine a skeptic who says the following: "Yes, I agree that it is far more reasonable to believe that you are reading a paper in Oviedo than that you are standing in the middle of a road; the first is far better justified than the second. More than this, when it comes to deciding what to believe, there are significant differences in the degrees of justification for various propositions about the external world, and in a large class of cases, it is entirely unproblematic as to what one should believe. I simply deny that the level of justification one reaches in any of these cases is sufficient to call the resulting beliefs cases of knowledge." This is not, of course, the skepticism of Descartes' First Meditation; it is, instead, a much more modest and less exciting form of skepticism. Let us call this view High Standards Skepticism. (Kornblith 2000: 26)

Though it's not absolutely clear to me, I take it that Kornblith's high standards skeptic is like my bold skeptic in that she makes the 'bold' (as I use the term) claim that we are speaking falsely whenever, even in ordinary conversation, we claim to 'know' things about the external world. That at least is how I'm inclined to understand what Kornblith means to be claiming when he says 'I simply deny that the level of justification one reaches in any of these cases is sufficient to call the resulting beliefs cases of knowledge.' If she is not in that way like my bold skeptic, her relevance to my project, and, in particular, why Kornblith thinks she is the type of skeptic that my solution provides an answer to, becomes completely opaque to me. And if the high-standards skeptic does make this bold claim, she strikes me as being in some ways I'll explain quite interesting—and also in some ways actually

bolder than the Descartes of the First Meditation, whom Kornblith does find quite interesting, but, I think, gets all wrong.²⁰

²⁰ Kornblith seems to be insinuating that Descartes is advancing something like Kornblith's 'full-blooded skepticism' in the First Meditation. He certainly thinks that Descartes is advancing a more strident skepticism than mere 'high standards skepticism'. Now, Descartes is at times frustratingly inexplicit in Meditation One about the force of the skepticism he intends to be advancing (in order to battle in the later Meditations). But I see no positive basis for thinking Descartes intends to be advancing anything nearly so wild-eyed as Kornblith's full-blooded skepticism anywhere there. The force of the skepticism that is supposed to be motivated by concerns about possible dreaming in paragraph five is particularly tough to discern. The key steps in this argument (premises, conclusions) are left largely inexplicit. But one claim that looks like a step in the argument is explicit: that 'there are no certain indications by which we may clearly distinguish wakefulness from sleep' (Descartes 1967: vol. 1, 146). Building on that clue, one guesses that Descartes is likely supposing some premise to the effect that if there are no 'certain indications' by which to make the distinction, and if perhaps certain other conditions are met, then any belief he might have to the effect that he is awake and not dreaming lacks a certain status, and then perhaps some other premise to the effect that if the above lack obtains, and perhaps if other conditions are met, then his beliefs in such ordinary things as that he is seated by a fireplace also lack a certain status—likely the same status. But what status could he have in mind? (What is the *force* of the intended skepticism?) Our best clue seems to be in the explicit step: What would 'certain indications' (or 'sure signs,' by other translations) be needed for? The natural answer would at least seem to be: in order to be *certain* or sure that one is awake, and then in order to be *certain* or sure of the likes of that one is sitting by the fire, etc. I suppose one could attribute to Descartes premises to the effect that (if perhaps certain conditions are met), one needs certain indications or sure signs that one is awake and not dreaming to have *any level of justification whatsoever* for thinking that one is awake, and then to have any justification whatsoever for thinking the likes of that one is sitting by the fire. But why? Why would *certain* indications or *sure* signs be needed for *that*? Would we be doing Descartes any favors by attributing such an argument to him? There are of course plenty of places to turn to for a more serious look at Descartes's dream argument than I have just given it, but my favorite is (Curley 1978: 46-69), which (very sensibly, so far as I can see) renders Descartes's conclusion as (emphasis added): 'None of my beliefs about ordinary-sized objects in my immediate vicinity are *certain*' (1978: 52)—which is a very far cry indeed from the skeptical conclusion that none of those beliefs have any degree of justification whatsoever. But here, though I see no basis for thinking full-blooded skepticism is being advanced, and though I certainly don't take Descartes to be advancing anything so wild, based just on Descartes's lack of explicitness, it also seems that even such a wild reading cannot be securely ruled out. But when we get to the skepticism advanced later in the First Meditation, seemingly wider in scope but even milder in force than the earlier dream-inspired skepticism, one based on skeptical possibilities that in Meditation Three Descartes calls 'very slight, and so to speak metaphysical' (Descartes 1967: vol. 1, 159), there do seem to be strong positive textual grounds for thinking Descartes never intended the skepticism being considered to be nearly as strong in force as full-blooded skepticism—and also good reason to think it isn't even meant to be as strident as high standards skepticism, where that includes the claim that ordinary claims to 'know' (or the close enough Latin equivalent) go false. Key hints to Descartes's intentions are to be found in explanations like the 'atheist geometer' passage in the Replies to Objections 2, where Descartes evaluates the beliefs of a geometer, who, remaining an atheist, has not followed Descartes's theistic escape from the skepticism established in Meditation One, and so who presumably is still subject to the skepticism advanced there. Descartes explains: 'That an atheist can know clearly that the three angles of a triangle are equal to two right angles, I do not deny, I merely affirm that, on

Sometimes we judge a claim to be ‘uninteresting’ because, however important it would be if true, we find the chance that it is true to be too slight. However, Kornblith’s charge against the high standards skeptic is the stronger one that *even if* her claim were right, that would *still* not be philosophically important news, as he indicates in this passage in which he likens his high standards skeptic to an eccentric Vermonter who is peculiarly demanding about what constitutes ‘cold’ weather:

What should we say to the High Standards Skeptic? On the one hand, I think it is quite natural to suggest that this skeptic has a deeply deviant view about the nature of knowledge, or at least, about how the term ‘knowledge’ should be used. But in practice, if confronted with such a skeptic, it would probably be wise simply to capitulate. “Let us use the term ‘knowledge’ as you do,” I would say. Nothing much hangs on this. Since the skeptic agrees that we can make meaningful and important distinctions about how well justified we are in various claims, and agrees with us about which claims we should believe and act upon, nothing much turns on it. It is 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. (Kornblith 2000)

This, and especially the very end of the above quotation, shows that in contrast to what we might call a ‘full-blooded Vermonter’, Kornblith is judging that his ‘high standards Vermonter’, and, by extension, his high standards skeptic, are making claims that would be unimportant *even if they were true*. (And after all, if Kornblith were thinking in terms that would make the incredibility of a claim suffice for it to be uninteresting, surely his full-blooded Vermonter, even though he is making

the other hand, such knowledge on his part cannot constitute true science, because no knowledge that can be rendered doubtful should be called science’ (Descartes 1967: vol. 2, 39). Important to our current concern is how positive Descartes’s evaluation of the atheist’s belief is, and how weak in force is skepticism that Descartes thinks the atheist is subject to. This seems weaker in force than even high-standards skepticism. I here use a translation of Descartes that is particularly friendly to the points I am making, but on any reasonable translation, there seems little room for the kind of understanding of Descartes Kornblith is pushing here. See (DeRose 1992b) for further discussion of how I understand Descartes on these matters.

‘substantive claims about the world’, would be among the very first in line to be written off as uninteresting!)

But that stronger judgment seems simply incredible to me. Perhaps this will end up just being a point on which Kornblith and I are extremely far apart in our thinking (and in ways I no doubt won’t be resolving here to the satisfaction of those who see things as Kornblith does, though I will say a bit about why I see things as I do), but this strikes me as implausible enough to warrant our looking for a not-so-obvious way of understanding Kornblith that makes his position more sensible than a surface reading would yield. And here I would note again that Kornblith’s paper is a response to (Sosa 2000), a main theme of which is an attempt to disentangle the elements of the contextualist position that concern epistemology from the aspects important to the philosophy of language. Perhaps Kornblith would agree that if our claims to ‘know’ things about the external world were false, that would indeed be news of considerable *philosophical* interest, but that interest would be confined to the philosophy of language, and the news would not be important to epistemology? After all, he does write things like (now with emphasis added): ‘The skeptical problem DeRose’s contextualism addresses is simply uninteresting *from an epistemological point of view*.’ Perhaps he’s thinking this news might just affect how some points of epistemology are expressed, but would leave the important epistemology itself unaffected, its considerable philosophical importance being confined to the philosophy of language? Or perhaps, on another reading, he thinks that what philosophical importance there is here is mostly confined to the philosophy of language, but he also thinks that importance is not considerable in any case?

At any rate, the news that all our claims to ‘know’ external facts, even those made in ordinary conversation, are false would certainly be startling—as indeed would be analogous news about ‘cold’. But as I think of how this astonishing news would affect my philosophical thinking, its *initial* impact would certainly seem to be the havoc it would at least threaten to wreak in the philosophy of language. One of the first things to do in response to such remarkable news would be to try to determine whether it was due to some isolated glitch affecting ‘know(s)’ (and/or ‘cold’), or whether common terms of our language more generally failed in such a radical way in allowing us to make true claims (and having true thoughts) about the world by their use. But especially as the threat tended toward the second of these possibilities, the news would seem far from ho-hum.

And I for one am skeptical that the resulting upheaval could in the end be confined to the philosophy of language. But this may well be based on ways that Kornblith and I just move to the beat of very different philosophical drums. *I* believe that separating out the factors which affect whether claims involving philosophically important terms are true or false from when such claims have other properties that might be confused with truth and falsity is an important tool for understanding the meanings of those key terms (and is a tool that would be largely disabled for use on a term once a high standards skepticism about that term were accepted), and that, in turn, *while it is certainly very far from all that is involved*, understanding the meaning of the key terms involved—or perhaps more precisely, avoiding radical and profound misunderstandings of those meanings—is important to addressing philosophical problems.²¹ If so, then to echo the thoughts of my Introduction to volume 1 (2009: 18-19), we would be bound to fall into error about knowledge, as well as about ‘know(s)’, by depriving ourselves of such important tools.

Before leaving the philosophical interest of Kornblith’s high standards skeptic, I should note the vital point that, while I’ve so far been addressing how important the news would be that this skeptic’s claims were right, there is also a non-hypothetical interest in this skeptic’s position, if that position includes not just her incredible claim, but also *arguments* for that conclusion, insofar as these arguments are strong and seem to powerfully support her bold claim that we are always speaking falsely when we claim to ‘know’ things. If this skeptic has a formally valid argument to the conclusion that we don’t really ‘know’ even in what seems to be a paradigmatic case of external world knowledge, and if that argument is driven by the premises that can seem boldly true, and especially if we do or are very tempted to judge that they are boldly true, then this is surely philosophically

²¹ I am sometimes mentioned as a current philosopher (such creatures were apparently more plentiful, or at least more prominent, in earlier times) who has ‘taken the linguistic turn’ in philosophy—sometimes as a particularly clear example of such. I’m not very confident in my understanding of this description, but I suppose that the sentence to which this note attaches would be the kind of spot to look at in deciding whether the description fits me. I’m guessing that my inclusion of the limiting phrase ‘while it is certainly very far from all that is involved’ may harm my candidacy. But if thinking that understanding the meaning of the key terms involved, and especially avoiding profound misunderstandings of the meanings of the key terms involved, can be important to addressing philosophical problems suffices for one to count as having taken the ‘turn’, then I certainly want to plead guilty.

important news, hopefully exposing various grave mistakes regarding knowledge that we are liable to make.²²

Similar points would hold for Kornblith's high standards Vermonter. However shocking and important would be the news that his claim is right, when thinking non-hypothetically, I'm inclined to write him off as a philosophically unimportant eccentric if he's just making his claim without any interesting support. But if we're talking about a high standards Vermonter *with a powerful argument for his bold claim* that we're speaking falsely whenever we say it's 'cold' outside and the temperature is above minus 25 degrees Fahrenheit, then that's someone that I, as a philosopher, want to talk to, even if his powerful argument isn't nearly powerful enough to rationally even tempt us to accept his conclusion. Though we will in that case have a right to suppose from the outset that his argument must go (badly) wrong *somewhere*, we'd still be likely to learn something important trying to figure out where and how it does go wrong.

11. Do I Respond Only to the High Standards Skeptic?: The High Standards Skeptic, the Deplorable Position Skeptic, and the Bold Skeptic

Charges like Kornblith's claim that 'DeRose responds only to the High Standards Skeptic' have been made by others, often taking the form that I, or that the contextualist responder to skepticism, fails to realize something very important about (some of) the most important skeptical challenges—namely, that the skeptic does not (in these important cases) think that we fail to know because the standards for knowledge are so high, but rather because we fail to satisfy even ordinary or perhaps even low standards. Thus, for example, here is Richard Feldman:

[Contextualists] fail to come to grips with some important arguments for skepticism. This is because in at least some cases the leading idea behind skepticism is not that we fail to satisfy some extraordinarily high standards for knowledge but rather that, contrary to common belief, we typically don't satisfy ordinary standards. This skeptical challenge gets remarkably

²² This is the value Greco sees in the study of skeptical arguments; see n. 19 of Chapter 2.

little attention from the contextualists under discussion here, since they seem simply to assume that we do meet ordinary standards for knowledge. (Feldman 2001: 78)

And, handily citing some of his predecessors, Kent Bach helpfully puts the charge (which he endorses) in a nutshell, impressively crystalizing the confusion:

Still, many philosophers have complained that contextualists do not really come to grips with the force and content of skeptical arguments (see Feldman, 1999, 2001; Klein, 2000; Kornblith, 2000; Sosa, 2000). The complaint is simple: the contextualist's attempt to marginalize skeptical arguments by restricting them to skepistemic contexts ignores the fact that skepticism denies that we have knowledge even by ordinary standards. The contextualist's attempt to marginalize these arguments by restricting them to skeptical contexts ignores the fact that skepticism denies that we have knowledge even by ordinary standards. (Bach 2005: 68)²³

This clearly represents some profound misunderstanding of what I at least am up to—and I'm fairly confident this would be so concerning other contextualists as well. For my intention, at least, is clearly to defeat, and not to ignore, the skeptic who 'denies that we have knowledge even by ordinary standards'.

Yet it is not hard to see what might be behind Feldman's claim that contextualists 'seem simply to assume that we do meet ordinary standards for knowledge': that does seem at least a fairly accurate observation of how actual contextualists, including myself, engage skepticism. So how could contextualists possibly be aiming to defeat a skeptic whose key claim is one that contextualists 'simply assume' is wrong? I can see how it can look like contextualists must be ignoring a skeptic who makes claims that we don't even meet ordinary standards for knowledge. And if contextualists are ignoring that skeptic, what kind of skeptic is left for contextualists to be really engaging? Perhaps just something like Kornblith's high standards skeptic? So I can see how this misreading of contextualists occurs.

We can start clearing up this muddle by noting the slippage between Kornblith's and my schemes for classifying skeptics. If Kornblith's high standards skeptic says that we're always speaking falsely when we claim to 'know' things

²³ Bredo Johnsen seems to have a charge like this in mind as well, here: '[C]ontextualists (along with many others) have badly misunderstood [the skeptic]; according to the skeptic, what accounts for my failure to know W is not that my belief lacks some extraordinarily high degree of epistemic strength, but that it has no epistemic strength' (Johnsen 2001: 394).

about the external world, which is the characteristic claim of my bold skeptic, then what is the relation between her and my bold skeptic? Answer: Kornblith's high standards skeptic is a more specific type of character: She *is* a bold skeptic, but there are other ways to be a bold skeptic, and some of them might be more interesting to Kornblith (and his fellow complainers).

Kornblith's high standards skeptic thinks our claims to 'know' facts about the external world are false because, when it comes to such facts, we never 'reach some very high standard required for knowledge'. But what of the different, complementary bold skeptic, so worrisome to Feldman & Co., who also thinks we never meet the truth conditions for ordinary claims to 'know' things about the external world, but thinks this not because those standards are so high, but rather because she thinks we're always in such a deplorable epistemic position with respect to such external matters that when it comes to them we don't meet even quite low standards for knowledge? Call this the 'deplorable position skeptic'.²⁴ This is the skeptic contextualists supposedly ignore, and she sounds like a kind of skeptic Kornblith at least is likely to find quite interesting.²⁵

So, how does this deplorable position skeptic compare with my bold skeptic? She too *is* a bold skeptic, every bit as much as is Kornblith's high standards skeptic. The deplorable position skeptic and the high standards skeptic represent two ways of being a bold skeptic, and when I defeat the bold skeptic, I mean to be defeating both of these characters. My *intention* certainly isn't to 'respond only to the High Standards Skeptic.'

²⁴ I realize that it's easy to feel that one is losing one's grip on the distinction between these two skeptics (high standards vs. deplorable position): How can the standards really be so low or moderate if we never, or almost never meet them? But I will not press that potential problem for my critics, instead showing that, insofar as we can make sense of the distinction between these two skeptics, I answer both of them.

²⁵ The deplorable position skeptic is much more like Kornblith's full-blooded skeptic, whom Kornblith thinks is of profound epistemological interest, than is the high standards skeptic. Kornblith presents his full-blooded skeptic as one who traffics in wild-eyed claims about (sometimes comparative) levels of justification—namely, that we have none of it at all for our external world beliefs, and no more for some such claims than for any others. The deplorable position skeptic we just met represents a way to fairly closely approach Kornblith's full-blooded skeptic while staying focused on knowledge rather than justification, and she seems to be a skeptic whom Kornblith would take to be making substantive epistemological claims, and rather strong ones at that, rather than just claims about language, thereby perhaps piquing his interest.

12. Do I Respond Only to the High Standards Skeptic?: The Simultaneous Defeat of the Bold Skeptics of Both Kinds

But how might such a defeat of two such different skeptics as the high standards skeptic and the deplorable position skeptic be accomplished in one fell swoop? And how can the contextualist's targets for battling include the second of those skeptics if the contextualist 'simply assumes' that this skeptic's key claim is wrong? A review of how the defeat is supposed to be accomplished will answer both of those questions, starting with the second of them.

Recall from the previous chapters the nature of the *defeat* that is being attempted, and its contrast with the other mode of combatting skepticism. The contextualist is not trying to *refute* any of these skeptics by deriving an anti-skeptical result from the perhaps extremely meagre stock of premises that do not beg the question against her. If *that* were the goal, then the way in which the contextualist 'simply assumes' that we meet ordinary standards for knowledge indeed would show that deplorable position skeptic could not legitimately be among the contextualist's targets. *But contextualist responses to skepticism have never been about that*—with respect to any skeptical targets. It is instead being deployed toward the very different (and philosophically more important, I claim) goal of providing the resolution of the puzzle that the skeptical argument presents us with that makes the most sense of all the intuitions involved, whether by endorsing them or by successfully explaining them away, and thereby defeating rival positions on the puzzle, including that of the bold skeptic. In that endeavor, it is to be expected that various questions will be begged against rivals—just as it is to be expected that those rivals, including the deplorable condition skeptic, will beg questions against the contextualist Moorean. We are supposing that advocates of the various positions on the puzzle have provided whatever positive arguments they can for the intuitive claims that they endorse, but have hit rock bottom, finding no still-more-powerful deeper claims on which to build, so we have reached the (hopefully) intuitive starting points of each position that will be 'simply assumed' rather than argued for, and we are now engaged in damage-control and explanation, and focused on the aspects of each position that are counter-intuitive.

And relative to the goal contextualists are pursuing, bold skeptics of both the 'high standards' and 'deplorable position' varieties are extremely vulnerable to such

a defeat, due to features of their position that, as bold skeptics, they hold in common.

The timid skeptic doesn't take AI to be showing that we are wrong when we say or think that we 'know' that O. Her characteristic claim is just that we don't 'know' by the elevated standards that the presentation of AI has a tendency to put into place, which seems far from a shocking claim, especially after we see all that contextualism about knowledge attributions has going for it (as we see in Volume 1). By contrast, the bold skeptic takes AI to be showing something very surprising indeed, to put it mildly. 'Shocking' probably isn't too strong a word to describe the bold skeptic's position that we have always been wrong when, in ordinary settings, we have said or thought that we 'know' this or that. And this is so whether our bold skeptic is a high standards skeptic or a deplorable condition skeptic—or something somehow in between those two.

In line with what I urge in section 10, the surprise of a bold skeptic's position, whichever variety of bold skeptic she is, portends well for the interest of her position. Shocking positions concerning such things as knowledge or truth are bound to be important news if true, and, even non-hypothetically, powerful arguments for such shocking positions are bound to be philosophically important.

But the shocking nature of the bold skeptic's claim, so helpful to her interest, constitutes a very serious challenge to her so far as the rational acceptability of her position goes, for reasons we also saw in Chapter 2. She is extremely vulnerable to defeat. For now she has to face the question of why anyone should rationally follow her reasoning to its bold skeptical conclusion, rather than accepting some other solution to the puzzle she confronts us with. For what does this bold AI skeptic have to stand on, other than how plausible her argument's premises are? As I've urged, the plausibility of her premises is considerable—and might even be more than just considerable. But the flip-side of the shocking nature of her conclusion is that its negation will also have at least a considerable degree of plausibility. And so our bold AI skeptics, of both the 'high standards' and 'deplorable position' varieties, face the very challenging question of *why in the world* it would be rational for anyone to accept her startlingly bold conclusion, rather than, say, concluding with the substantive Moorean, that, at least by the standards that govern our ordinary thought and talk, we really do 'know' that not-H. This Moorean conclusion may constitute something of an unexpected intellectual surprise for some, as our bold skeptics might do very well to point out, but many might quite rationally find it

much less incredible than the arresting shocker the bold skeptics are asking them to swallow. In short, as we have seen, what we face here is a puzzle, and the incredible nature of bold skepticism is what renders it, in all its varieties, highly vulnerable to defeat by other solutions that might come along.

So, when a contextualist proposes a solution that he claims really does *solve* the problem, by accounting for the plausibility of all the claims constitutive of the puzzle (rather than just embracing one of the initially distasteful claims), and argues that his solution is superior to those put forward in defense of bold skepticism, and claims to have thereby rendered it rationally preferable to accept his solution to the puzzle over that of the bold skeptic, it is wrong to charge that he has only responded to one sub-type of bold skeptic. The charge that the contextualist is not answering the deplorable condition skeptic *could* make sense *if* the deplorable condition skeptic were somehow immune to the problems the contextualist was pressing against bold skepticism, and/or if the complainer were putting forward as promising some proposed real solution to our puzzle that was especially available to the deplorable condition type of bold skeptic. But those making the charge don't seem to have anything like that in mind. Or, at least if they do, they are not very forthcoming about why they think the problems being pressed against bold skepticism don't affect the deplorable condition sub-type of the view, nor about what solution to the puzzle, especially well-suited to the deplorable condition variety of bold skepticism, they think is promising and is being ignored.

It is true that when in SSP I explore possible solutions available to the bold skeptic, I end up trying out a solution the adopting of which would tend to make the bold skeptic a high standards, rather than a deplorable condition, skeptic, and this may have caused some of the confusion here.²⁶ But that's not because the deplorable condition skeptic has a good solution up her sleeve that I'm aware of but am choosing to ignore. Rather, the relevant part of SSP (sections 15-16) unfolds as follows. I note the problems that the bold skeptic faces, and, though I don't break things down into cases (which would have been needlessly tedious), the problems I press affect bold skeptics of both varieties. I then note a kind of maneuver by which it has been suggested that bold skeptics might deal with their problems, and actually

²⁶ My suspicion, however, is that the confusion was instead largely produced by the thought that we've already addressed that the contextualist must be ignoring the deplorable condition skeptic because the contextualist 'simply assumes' that such a skeptic is wrong.

solve the AI puzzle. I then note a couple of quite fundamental problems with that maneuver, and, again, the problems I allege wouldn't be confined to the use of the maneuver by just one or the other sub-type of bold skeptic, but would affect both of them. I then note a way that a bold skeptic might try to deal with her new problems, and it is at this point that the possible escape I consider on the bold skeptic's behalf would seem to push her toward being a high standards, rather than a deplorable condition, bold skeptic, I think. I then argue that my contextualist solution proves a better solution than the one I have been led to consider on the bold skeptic's behalf—and again, the problem here, though it's at this point being pressed against what looks most like a high standards skeptic, is one that afflicts both varieties of bold skeptic. So, at the close of this story, I end up dealing with a bold skeptic who is looking much like a high standards skeptic. But that doesn't mean the deplorable condition skeptic isn't being targeted. She is still among the rivals who are supposed to be defeated. I have said why I think the bold skeptic's solution fails, and the problems I base that conclusion on, up to the end, are problems that afflict the deplorable condition skeptic in her attempt (if she even makes the attempt) to solve the AI problem as well as the high standards skeptic. The reason the deplorable condition skeptic gets neglected a little bit at the end of the story is because it is her fellow bold skeptic, and not her, who seems to me to have a potentially promising move yet to try.

Of course, it is possible that there are promising damage-controlling moves available to the bold skeptic, and perhaps more specifically to the deplorable condition skeptic, that I haven't heard of and just can't see. If that's the complaint, then the way to press it would be to present those moves. As I admit in SSP, 'Like practically any claim to have provided the best explanation of something, my claim here is hostage to the possible future development of a better explanation coming along.' But if I am missing some promising possibility for the deplorable condition skeptic, I'm not doing so knowingly; she is at least not being intentionally ignored. She is among the rivals who are supposed to be defeated in the way I have explained.

In fact, some might well object to my claim to have defeated the bold skeptics of either variety on grounds quite unfriendly indeed to those skeptics: that bold skepticism was already pre-defeated independent of my anti-skeptical efforts—or those of any other anti-skeptics, for that matter. The sheer incredibility of the position has defeated it, far above any epistemologist's poor power to add or

detract. This will be the stance of those who make their ‘Moorean choice’ quite definitively against the bold skeptic, finding her solution far inferior to some other on offer (e.g., bare, bullet-biting Mooreanism), and who aren’t impressed enough by the potential of the kinds of damage-control maneuvers available to bold skepticism to begin to even begin to tempt them to, on their basis, make an ‘enlightened Moorean choice’ that’s any different from, or more difficult for them to make than, their original, unenlightened choice. For them, the bold skeptic did not need to be defeated by the development of some other solution and/or any arguments about the viability of damage-control maneuvers. Still, they take the bold skeptic to be defeated, and they may still be in the market for an explanation of how this puzzle arises, even though the stance they take toward bold skepticism, at least insofar as it’s motivated by AI, does not depend on any understanding that such an explanation might provide. If one with such strident anti-skeptical leanings were to find a solution like mine to be both successful and the best on offer, I imagine they could see their way clear toward applying a perhaps somewhat extended use of our notion of ‘defeat’ to what the ‘winning’ solution does to rival positions, including here that of the bold skeptic, even if those solutions were already defeated (in a stricter sense) before the solution in question was developed.

13. My Supposedly Thoroughly Externalist Response to the Full-Blooded Skeptic

Though Kornblith complains about my ignoring his full-blooded skeptics, at one point he says I have a response, but claims that what I response I have has nothing to do with contextualism:

DeRose does have a response to Full-Blooded Skepticism, but it does not involve his contextualism. DeRose, in addition to being a contextualist, is an externalist about justification. And DeRose’s treatment of what he calls the “strength of one’s epistemic position” is a familiar externalist account. On such a view, the skeptic is simply mistaken when he claims that the conditions for knowledge cannot ever be fulfilled; he is mistaken in thinking that in a situation of the sort we ordinarily take ourselves to be in when looking at a table (in standard conditions with properly functioning perceptual and cognitive equipment), we are no better justified in believing that we are looking at a table than that we are looking at a toaster. The externalist account of justification, or strength of epistemic position, is a

familiar and controversial one. It may or may not be right. I myself am very sympathetic with externalism. But it is important to recognize that it is the externalist part of DeRose's view which is doing the work in combatting Full-Blooded Skepticism. Contextualism does no work here. (2000: 27)

And one encounters similar thoughts from others that it is really externalism (and a crude form of it) that does the work for me.

First, a small point: Before this book, I have not (at least in print) addressed skepticism about justification, which Kornblith's full-blooded skeptic advances, in a way that involves contextualism. (I will address questions of justification in Chapter 7 of this volume.) But it turns out that I'm an internalist about epistemic justification—not that Kornblith should have been able to tell that (though I don't know why he thought he could tell I'm an externalist about that). However, as we've now seen, I do take myself to be responding to a skeptic not that far from Kornblith's full-blooded skeptic: our deplorable condition skeptic, who, as I've noted (note 25), represents a way to closely approach Kornblith's full-blooded skeptic while staying focused on knowledge rather than justification.

And here we come to the important point: I haven't responded to any skeptic in the way Kornblith alleges above. Though I stress that it is 'rough and ready', I do use an account of knowledge in constructing my explanation of how the AI puzzle is generated, and then in my attempt to solve that puzzle. And this account, which I've since come to call a 'double-safety' account, does articulate what can well be thought of as an 'externalist' condition for knowledge. Hence, it would be an externalist account of knowledge if it, together (one supposes) with the standard conditions on knowledge (true belief) were taken as a complete theory of knowledge. And it would certainly make for a (*very implausibly, to my sensibilities*) externalist account of justification. If I had simply put forward such an account of knowledge or justification, and then rejected some form of skepticism about knowledge or justification because such skepticism was inconsistent with my theory of knowledge or justification, together, I suppose, with my view of what the world is like (to secure the result that worlds where I go wrong with respect to O really are all quite distant), then Kornblith's reaction would be on-target. But that's not what I've been up to. (That's to take me to be engaging with skepticism more in the way that Nozick

does.²⁷) My response to bold skeptics has *not* been: ‘Here’s my complete (and thoroughly externalist) theory of knowledge; and look, given my view of what the world is like, we do meet every condition for knowing that O, at least by ordinary standards for knowledge, that my theory posits; so, you’re wrong.’ My focus has not in that way been on explaining how we know, but rather on why it can seem that we don’t. So my response to the bold skeptics is better thought of as being along these lines: ‘On my view, of course, you’re wrong; as I am on your view. But I can on my view explain (away) the source of attraction to your view, like this What do you got?’ It’s in the task (crucial to the project of *defeating* the bold skeptics, as we’re using that term) of explaining (away) the intuitive power of the skeptic’s argument that my account of knowledge is utilized. Its role then does not require it to be a complete or finished theory of knowledge (thank goodness!), nor even to be close to a complete theory, so that our satisfying it would at least provide good reason to suppose we probably do know, but just to provide the materials for the needed undermining explanation. My use of the double-safety account doesn’t then commit me to a radically externalist view of knowledge, but just to an externalist component of knowledge being able to play the crucial role my solution assigns to it in accounting for the intuitive power of the skeptic’s claims.²⁸

Moving now to respond directly to Kornblith’s key charge: Once we are clear on how I respond to bold skeptics, we can see how wrong it is to suppose that it’s some externalist account of knowledge that provides my response to any skeptics in a way that does not involve my contextualism. For the account’s only role in my response is to provide an explanation of how the skeptic’s argument can make it seem (to some extent) that we don’t know—to explain (away) the intuitive power of the skeptic’s case. And it cannot play that role in anything like the way I employ it without the contextualism. It’s my contextualism and my (partial and rough-and-

²⁷ See esp. the first two sections of Chapter 7 for a more sustained explanation of the difference between what I there characterize as Nozick’s positive explanatory approach to skepticism and the negative explanatory approach I am attempting in the work that Kornblith is critiquing. See the rest of Chapter 7 for my own attempt to now add a positive explanation to my treatment of skepticism.

²⁸ The supposedly externalist nature of my theory of knowledge is a main theme of Bredo Johnsen’s account of my treatment of skepticism at (Johnsen 2001: 392-4). Though I do not here directly respond to Johnsen, I hope that my response to Kornblith in this section makes clear how I would answer Johnsen’s related concerns. As it happens, before either paper was published, I encountered Johnsen’s work before Kornblith’s.

ready, but, I claim, enough to get the job done) account of knowledge working together that provide the crucial explanation. Neither works alone.

Of course, one could wonder how well the rough-and-ready account of knowledge I use could be advanced into a more complete account of knowledge and made to fit with a plausible account of justification while still allowing for the kind of explanation I give in response to skepticism. These questions are taken up in Chapter 7, where I now provide the kind of positive anti-skeptical explanation of how we know what we claim to know that I have been misunderstood as giving in my past work (mainly, SSP).

14. Indexed AI

We now turn to important criticisms of the contextualist strategy leveled by Sosa. We start with a challenge based on a modified version of AI that Sosa puts forward (Sosa 2004a: 50-51). Using the subscript 'o' so that 'knows_o' designates knowing by the standards for knowledge that tend to govern ordinary, non-philosophical uses of the term, the argument is just like AI, except all occurrences of 'know' are explicitly specified as being governed by o:

Indexed AI

- 1i. I don't know_o that not-H.
 - 2i. If I don't know_o that not-H, then I don't know_o that O.
- So, Ci. I don't know_o that O.

Sosa thinks this argument does about as well as the original in motivating some of the skeptic's most important worries, but is immune to the contextualist's response because in this version of the argument, the 'contextual terms have been relativized away' (2004a: 51).

The contextualist Moorean (like me) will hold that Indexed AI is unsound, and that it goes wrong at its first premise, which is simply false. The challenge for us would be to account for the plausibility of that premise. Though, for the reason given in note 31, I don't think 1i is as intuitively attractive as 1, it is nonetheless quite plausible, I'm inclined to agree, so this is a real challenge.

But it is not really a new challenge. At least for me, this was the crucial part of how I viewed the challenge I initially set out to meet, and the ability of the contextualist's response to account for just what's at issue here—how AI's premises, and, here, focusing on its first premise, can seem to be true by ordinary standards for knowledge—grounded the importance of the contextualist strategy in the first place. As I wrote in section 2 of SSP:

The contextualist strategy is important because AI initially seems to threaten the truth of our ordinary claims—it threatens to boldly show that we've been wrong all along in thinking and saying that we know this and that. For it doesn't seem as if it's just in some "high" or "philosophical" sense that AI's premises are true: They seem true in the ordinary sense of 'know'. In fact, one is initially tempted to say that there's *no* good sense in which I know that I'm not a BIV or in which I can know I have hands if I don't know that I'm not a BIV.

On my account, it's the insensitivity of any belief we might have that not-H that accounts for the plausibility of the skeptic's denial that we know that not-H. Though I do not take sensitivity to be a requirement for knowledge (which becomes important when AI's second premise is addressed), on the skeptic-friendly version of contextualism I provisionally assumed in SSP, this insensitivity secures, via the 'Rule of Sensitivity', the result that any attempt to claim to 'know' that not-H (or any attribution of 'knowledge' of not-H to some person other than the speaker) will be false, and any attempt to admit that one does not 'know' that not-H (or that some person other than the speaker does not 'know' that not-H) will be true. This account of the plausibility of the premise would not lead us to expect that the premise will be attractive to us only when we are understanding it to be invoking some special, philosophical interpretation of 'know(s)'; after all, it predicts that *any* claim to 'know' that not-H, in any kind of conversational setting, will go false (and likewise, any admission that one doesn't 'know' that not-H will come out true). If we take the use of ordinary standards to be necessary to a context's counting as 'ordinary', this account implies that no context in which such a claim (or admission) is made will remain ordinary. Still, it implies that in any context in which the claim is made, no matter how mundane (at least up to the point when the claim is made), making the claim backfires. Insofar as our judgment about the truth value of a claim is influenced by our sense of whether it would be true or false to actually make the claim, this account seems well-suited for explaining even why it can seem²⁹ that

²⁹ At least to some extent: the appearance here has different strengths for different people; see sections 7-10 of Chapter 2 for discussion.

there is *no* good sense in which we know that not-H, since it predicts that there is no context in which one can truthfully make the claim.

The above, though, is just my account under the provisional assumption of a skeptic-friendly version of contextualism. On a more serious (at least by my lights) version of contextualism, it won't predict that just any claim to 'know' that not-H will be false. However, it will still imply that any such claim will have a force working against its truth—a force which, while it may on occasion be overpowered by some other conversational force, would result in the claim's falsehood were it unopposed. A claim to know not-H (where one's belief that not-H is insensitive) will then likely never seem unproblematic. What's more, on some of the less-skeptic-friendly views, including the one I'm inclined toward (DeRose 2009: 144-8), all it would take to make a claim to 'know' that not-H to go non-true (even if not false) is some skeptic stating that you 'don't know' that not-H. It would be easy for us to write off any sense we might have that such a claim has some truth-like virtue as being illusory. After all, what kind of truth disappears upon meeting verbal resistance? Of course, we now know that on contextualism, nothing that can properly be called 'knowledge' appears or vanishes, depending on what kind of reception one's claim to 'know' meets (unless the reception one's claim meets rocks one's belief in the proposition in question): Regardless of context, one knows_o that not-H and does not know_h that not-H. It is only the ability to express a true proposition by means of a given sentence, and not the truth-value of any given proposition that might be so expressed, that is really in jeopardy here (DeRose 2009: Chapter 6). Still, prior to all that theory and clarification, we can be expected to look askance at claims whose truth requires a certain reception (or at least the lack of a hostile reception) by interlocutors.³⁰ Since my account has it that there is always some important force working against the truth of any claim to 'know' that not-H (where H is well-chosen), and since efforts to resist this force will likely strike us as at least somewhat suspicious, it is the reverse of surprising that we will quite generally find the truth of 'I know that not-H' to be at least problematic. Thus it is no surprise on our account of the plausibility of the skeptic's first premise that we

³⁰ This suspicion can and I think does fly in both directions (at least on different occasions). Those with a strong sense of the truth of their ordinary claims to know that O (those with certain strong (substantively) Moorean tendencies) may find the skeptic's denials deeply suspicious in large part because those denials can be true only in contexts meeting certain conditions.

would have at least some tendency to judge that that premise is true even by ordinary standards for knowledge, or even by *any* good standards for knowledge.

Thus, the contextualist solution to AI can also account for Indexed AI. In fact, it was all along in an important way an attempt to account for the former in a way that could also account for the power of the latter.³¹

15. Irrelevant to Traditional Epistemological Reflection on Skepticism?

We now turn to Sosa's worry that contextualist solutions to skepticism, including mine, are simply irrelevant to traditional epistemological reflection on skepticism.

It turns out that it's only on a very narrow construal of what traditional reflection on skepticism has encompassed that this worry is realized—an implausibly narrow construal, if you ask me, but if it were accurate, that would just mean that traditional reflection wasn't all that valuable.

Sosa compares

(a1) People often utter truths when they say "Somebody loves me."

(a2) Does anybody love me?

with

³¹ Of course, this isn't to say that the two arguments have the same amount of intuitive power, generally or specifically at their first premises, and I should note that I don't think that responders will generally find 1i to be as plausible as they would have found 1 to be. Indexed AI, already at its first premise, will raise the possibility of contextualism in a way that the original skeptical argument does not. One presented with the original version of AI might not consider the possibility of contextualism, and so might be strongly inclined to presume that they either simply know that not-H, or simply don't know that. Such a responder may be more easily coaxed into accepting 1 than 1i, but may still be taking 1 to be boldly true, in that if they judge 2 and then C also to be true in the same spirit, they will take this to imply that they've been wrong all along in thinking and saying that they 'know' the likes of O. By contrast, if such responder is presented instead with 1i, this will tend to raise the possibility of (at least something like) contextualism as a consideration. This may make them somewhat less inclined to accept 1i than they would be to accept 1.

(c1) People often utter truths when they say, “I know there are hands.”

(c2) Do people ever know that there are hands?,

adding: ‘c2 is presented as a question we might pose in philosophical reflection, in a philosophy journal or conference.’ A distraught person led to ask a2 is hoping to be reassured by the answer that indeed somebody does love them or, failing that, is looking to have their worst fears to the contrary confirmed; a1 miserably fails to provide the comfort hoped for. In philosophical reflection on skepticism, in philosophy journals, and in philosophy conferences, many have been moved to ask (the likes of) c2. Why is c1 any more relevant to c2, asked in a philosophical setting, than a1 is to a2?

Let’s for the time being make the ‘skeptic-friendly’ assumption that philosophical discussions of skepticism are governed by exceedingly high epistemic standards according to which nobody knows that there are hands. (We can call these ‘absolute’ standards.) As you now know, I’m in the end inclined to reject that assumption, but making the assumption helps to make the question of relevance more pointed, and you by now know how to convert what we say while working under such an assumption to what I really think. Under that assumption, when c2 is asked in philosophical discussions of skepticism, the truthful answer to it, according to contextualist solutions to skepticism, is the potentially distressing ‘no’. The pressing question: Why is c1 any comfort here — any more than a1 is in the distressing situation in which the answer to a2 is the distressing ‘no’?

A crucial difference between a1/a2 and c1/c2 is that the relevant ‘context-sensitivity’ in a1/a2 — that ‘me’ refers on each occasion of use to the speaker, and so refers to different people as it is spoken by different people — is clear enough to all that nobody (well, except for the likes of young children, who can be caught up in various, often quite cute, ‘paradoxes’ here) is likely to be misled into thinking that a1 implies that the answer to a2 is positive. By contrast, if the contextualist treatment of skepticism is correct, it is only controversially so, and it is far from being clear to all.³² Under contextualist analysis (and for the time being, a skeptic-friendly version of it), the skeptic raises the standards for knowledge, and our sensitivity to her

³² If you think the controversy and lack of clarity here themselves favor invariantism over contextualism, note that it’s also very much the case that if contextualism and contextualist treatments of skepticism are *wrong*, they are only controversially so, and recall the argument of (DeRose 2009: 159-60, 177-79).

standards-raising maneuvers inclines us to at least some extent to give negative answers to questions about whether there is knowledge in various cases. At the same time, however, we are sensitive to the fact that it's usually correct to say that people do know all sorts of things — i.e., to some extent, we are still inclined to realize that c1 . Not realizing that those usual affirmations are compatible with skeptical denials of knowledge, we naturally get confused.

So far, that all sounds very skeptical! (Though this is in part due to our working under the skeptic-friendly assumption.) The correct answer to c2, as we've been asking it in philosophical discussions, is 'no', and our tendency to think otherwise is based on a confusion — the confusion of thinking that c1 implies that c2 should receive a positive answer in philosophical settings.

But crucially, the confusion, of course, pushes both ways. Not only does c1 push us toward thinking that c2 should be answered positively, but we can also easily think that a negative answer to c2 threatens to imply that c1 is really false. And that, I think, is a large part of why a negative answer to c2 (made during philosophical discussion) can seem so menacing and so important—or perhaps even why that answer seems significantly important at all. I mean, really, doesn't at least much of the sensed threat and interest of skepticism come from the thought that *perhaps we've been wrong* all along in thinking and saying that we 'know' various things? If it instead turns out that almost all our thoughts and assertions to the effect that we 'know' various things — including very serious thoughts and assertions made in very serious settings where appropriately high standards (but not yet philosophical, absolute standards) for knowledge hold sway — are really correct, and especially if it turns out that the skeptic never even meant to be suggesting otherwise, skepticism loses most of its sting, and interest, at least for me.³³ Therein lies the comfort of the contextualist response.

³³ Compare this from Cohen, whose attitude I am here at least fairly closely following: 'What is truly startling about skepticism, is the claim that all along, in our day to day lives, when we have claimed to know things, we have been wrong—we have been expressing propositions that are literally false. If the skeptic's position is interesting, it is because he challenges our everyday knowledge attributions. The theory of relevant alternatives addresses the challenge by arguing that our everyday knowledge attributions—properly interpreted—are correct. The propositions we actually express and have been expressing all along are literally true. We do know relative to the standards that ordinarily govern those attributions. Now it is of considerably less interest (although not devoid of interest) that there is some other proposition involving

Is that sufficient comfort? How much skeptical distress remains unaddressed? That depends on how important it is that we know according to the 'absolute' standards. In an extreme case, if the hazard to our absolute knowing is the only threat of skepticism that one finds at all important, one will derive no comfort at all from the contextualist solution. I of course don't have that extreme view and see this very differently. I have elsewhere explained what I think it would take to meet the absolute standards in question, and on my proposal, it is simply impossible for any conceivable being to meet those standards for any belief it might hold (DeRose 1990: 292-309). But as my account of that will be controversial, I should register here that, quite independently of my particular, controversial account of what it would take to meet the absolute standards, for my part, once the skeptical strategy is seen to have no tendency to show that any of my claims to 'know' — except those very rare ones made in settings governed by 'absolute' standards — are in any way wrong, and once I start to get a look at what it *would* take to 'know' according to the skeptic's absolute standards, on *any* reasonable account of that, I find the distress caused by my failure to meet those standards to be very slight, and so to speak metaphysical (to put things in properly Cartesian terms)—perhaps to be compared with the 'distress' produced by the realization that I'm not omnipotent. (And if my above-mentioned suspicion is correct, it is perhaps well compared with the distress caused by the realization that I'm not 'omnipotent' on a construal of that attribute on which it is incoherent to suppose any conceivable being could possess it).

c1 does not tell us how c2, interpreted according to absolute standards, should be answered—any more than a1 tells one how to answer a2. If traditional epistemological reflection on skepticism were strictly limited to the narrow question of how to answer c2, when it's interpreted in the absolute way in which it seems to be used only in quite rare discussions of skepticism, if even there (we are still to decide whether our skeptic-friendly assumption is right), then Sosa's worry would be valid: The contribution of the contextualist solution to that project would be to tell us that the skeptic's answer to the lone important question of traditional inquiry about skepticism is the right one and that any tendency we might have to reject the skeptic's answer that's based on facts like c1 is misguided. Of course, the

stricter standards that is false. It is not the case that we know relative to skeptical standards. This a fallibilist readily concedes' (Cohen 1988: 117).

contextualist solution might very well also, in that case, lead us to conclude that traditional inquiry into skepticism has not been concerned with a very important question. I would certainly be inclined to judge so.

But surely traditional philosophical inquiry into skepticism has been concerned with more than that unimportant question! Insofar as philosophical inquiry into skepticism concerns, not just the question of how c2, interpreted absolutely, should be answered, but also concerns the importance of that question; insofar as it addresses the relation between philosophical denials of 'knowledge' and ordinary thought about 'knowledge'; and, crucially, insofar as philosophical inquiry into skepticism seeks to address the truly menacing thought that c1 might not be correct after all, the contextualist solution is far indeed from irrelevant. And insofar as there's room in traditional philosophical inquiry into skepticism for that menacing thought to be addressed by an investigation of whether the skeptic's argument renders it reasonable to accept that disturbing thought, then, as I would have thought SSP has made evident all along, but as I certainly hope these past three chapters have made clear, that's what my contextualist treatment is all about. *I* think that is just the kind of question traditional approaches to skepticism have been interested in. But if I'm wrong about that, then I say: so much the worse for traditional inquiry into skepticism.

16. Is My Contextualist Position Inexpressible?

We now move into objections the evaluation of which involves the mechanics of just when and how the epistemic standards for knowledge claims are moved. We start with charges that my contextualist position on skepticism is inexpressible.

This is for me a very old objection, predating any actual writing up of my position, even in the draftiest of drafts, as it was raised by my advisor when I was just talking through my basic approach to AI that was at the time just starting to make sense to me. Here's how I recalled (as best I could) the event on Facebook:

A quick recollection of working with my dissertation supervisor, Rogers Albritton:

When I was a grad student, and Albritton had bad problems with his sciatic nerve, I had to take over teaching his upper division undergrad class on skepticism. So I had to drive over

to his condo a few times, ostensibly to talk about what to cover in class and how, but (as anyone who knew him could have predicted), it turned out just to be a discussion of the issues themselves, with basically zero attention to issues about how anything should be covered in class. He was laid up in bed, and I'd sit in a chair across the room. One time, the discussion went very long (which also could have been very easily predicted), and, since it started with the bright afternoon sun shining through the window, with no lights turned on, it ended in total darkness, because we never bothered to turn on any lights. So, I ended up with the experience of his voice coming at me from out of the darkness, like the voice of God. (And if I had to cast someone as the voice of God, he'd be one of my first choices. Him or Orson Welles, depending on which personality traits of God's were being highlighted by the role.) And I still remember just what he said to me at one point, out of the darkness, and how he said it, with the dramatic pause before the last word, which was delivered as only he could, dripping with playful mock disdain:

“So, according to you, there are these facts to the effect that, at least by ordinary standards, we know various skeptical hypotheses are false. But we can't state these facts, because the very attempt to do so drives the standards up to where we no longer know. So if we can't speak them, what are these facts *doing*, just ... obtaining???!”

My response: “They're facts. What'dya want them to do?”

Which got the biggest laugh from him I think I ever got.

There was of course something odd about Albritton's presentation of the alleged problem that on my view 'we can't state these facts' (and this oddity has also characterized the presentation of many who have confronted me with this problem in the subsequent years): He had little problem (following the wording of my own earlier attempts to do so) stating just what aspect of my view was supposedly unstatable. At least if I'm right about how phrases like 'by ordinary standards' can work, I can employ a little 'semantic ascent' and express the relevant thought in the way the critic just did (usually, following my verbal lead): 'By ordinary standards, I (or we) know that H is false.' (And if I'm wrong about how 'by ordinary standards' works, we could certainly find some other, though perhaps more ugly, way to engage in a little ascent and say what needs to be said here.³⁴) So perhaps the objection should be refined to the charge that on my view, one cannot express the

³⁴ See (DeRose 2009: 215-17) for examples of 'easy little phrases' that can be used (at least according to me) as handy devices of semantic ascent and some examples of how they can be so used, but then also for a little guidance on how to resort to more 'tiresome' ways of achieving that effect if you don't go for my use of the handy little phrases.

facts I allege *by means of very simple sentences* of the form ‘I know that not-H’. I’ve since heard such inexpressibility objections, whether refined or not, many times, and read it a few times.³⁵

I am quite uncertain about whether and why the refined charge should be troubling, especially once we distinguish it from the potentially more-troubling-if-correct, but clearly wrong, unrefined charge. So long as the relevant part of my view can be easily expressed by means of fairly simple sentences that involve a little ascent, would it really be troubling if it couldn’t be expressed by even simpler sentences? Sometimes, the truth of philosophically puzzling matters isn’t maximally simple. We should have been thankful that this one is as simply expressible as it is.

But, however untroubling it would be if true, even the refined charge is false, in any case. The reason it can be perhaps somewhat *difficult* to truthfully state ‘I know that not-H’ on my view is that there is a conversational rule—my ‘Rule of Sensitivity’—by the operation of which, when such a claim is made, the epistemic standards *tend to* be raised to a level at which the claim is false. But that’s just a tendency, and the operation of that rule can of course run up against, and lose out to, other conversational forces that can work to keep the epistemic standards meetably low. And indeed, I gave some examples of what I thought, and think, to be true claims to ‘know’ that skeptical hypotheses are false (despite the fact that the ‘knower’s’ beliefs that they’re false are insensitive)—cases that appear in Chapter 6, section 8 of this book.³⁶

³⁵ (Sosa 1999: 144-5), (Koethe 2005: 76), (Davis 2004: 206).

³⁶ In fairness to those who have misunderstood me here, I did not in SSP itself actually present the cases of truthful attributions of ‘knowledge’ to insensitive believers that I had devised, but only mention in a footnote that there are such cases, referring readers to my (already by then, several year old) dissertation (DeRose 1990) for the actual cases (which are now found in Chapter 6 of this volume). What these critics were missing, in addition to that note, is the second of three ‘important points’ I make in setting up my contextualist solution in section 2 of SSP, where I explain that I will be provisionally assuming, for ease of exposition, a ‘skeptical-friendly’ version of contextualism, on which the raising of skeptical hypotheses always succeeds at installing the relevant elevated epistemic standards, and then explain how to convert my solution to less skeptical-friendly versions of contextualism. Because I was making that working assumption, much of SSP does read as if I hold the view that these critics misattribute to me. It was in (DeRose 2004a) that I started to explain how ‘skeptical-(un)friendly’ my own contextualist thinking really is. And such misunderstandings of my view do seem to have become less frequent since then.

And, as we will see toward the end of the next section below, there is another way to see to it that one's simple claims of the form 'I know that P' are governed by moderate epistemic standards (on which one does know that not-H, on my view): explicitly state that one's uses of 'know' are to be so understood—as I will have occasion to do myself in this book, in (the fourth paragraph of) Appendix C. Having done that, one can, in my opinion, truthfully assert the very simple sentence 'I know that I'm not a BIV'.

In case it helps, I can also do just that right now, too, producing an instance of the type of claim in question that *I* at least hold is true (though I admit its truth-value is far from obvious). The following claim that closes this section should be understood as governed by ordinary, moderate standards for knowledge.

I know that I'm not a brain in a vat.³⁷

Davis seems to catch his mistake in a footnote: 'DeRose's (1995, p. 39) most official formulation of the rule says that epistemic standards "tend to be raised, if need be" so as to require sensitive belief. But then his theory does not completely avoid abominable conjunctions, and needs to explain when shifts do occur' (Davis 2004: 278, n. 6). But if Davis does indeed see that I hold that it is possible to truthfully assert 'I know that not-H', he for some reason retains, in his main text, the charge that on my view this is impossible, so I am unsure whether he means to be withdrawing the charge (replacing it with new charges). He does not consider my cases where I hold that such claims actually are true, my attempts to explain why that occurs, nor any of my attempts to explain when shifts do and do not occur. As for why conjunctions can be 'abominable' (or can 'clash') even though the two conjuncts can each individually be given readings on which they are true, see (DeRose 1998: 70-2)—though this doesn't go all that far beyond the fairly obvious observation that conjunctions involving context-sensitive terms often do 'clash' (give an appearance of inconsistency) when the 'conversational score' is held constant over the whole conjunction, even though, if you allow yourself to change the score mid-sentence, you could assign to each conjunct a reading which would make the whole conjunction come out true. This is why the feature of my view I appeal to in section 10 of SSP to explain the abominableness of abominable conjunctions—namely, that they are false when evaluated consistently at any epistemic standard—does do a good job of explaining what it seeks to explain.

³⁷ At least as I'm inclined to construe the position, it may be open to the classical invariantist, and even the skeptical classical invariantist, to hold that I've just asserted a truth in the sentence to which this note attaches. As I am inclined to construe the position, it concerns 'normal' claims about 'knowledge', in some good sense, where claims whose meaning is explicitly stipulated as I did in the text are not in that way 'normal', but are stipulated, technical uses, very different from ordinary uses of the term. (It isn't easy to see how my stipulation that I am using 'ordinary, moderate standards for knowledge' will be treated by the views in question. The skeptical invariantist holds that ordinary claims are governed by extremely high standards, so they could hold that my stipulation invokes those standards—in which case my claim is false. Or, so far as I can see, they could think my stipulation invokes the moderate standards claims must meet before we will in make them in ordinary circumstances, in which case they may think my claim is true.) From my contextualist standpoint, by contrast, the epistemic standards that govern knowledge claims are always moving around,

17. The Factivity Problem

Using the subscript 'h' so that 'know_h' designates knowing by the high standards that the presentation of AI has some tendency to put into place, and 'o' so that 'know_o' designates knowing by the ordinary epistemic standards that govern everyday talk (and we're likely idealizing here, but harmlessly so, in supposing there is a single such standard so described), claims along the lines of both of the following are central to contextualist treatments of skepticism, including mine:

- a. I know_o that I have hands
- b. I don't know_h that I have hands

After all, a central aim of contextualist treatments is to protect the likes of (a), and appeals to the likes of (b) play a key role in our protection efforts. But since knowing something by any standard, including o, requires the truth of the proposition known, it seems that since we contextualists accept (b), we should also accept the likes of

- c. I don't know_h that I know_o that I have hands.³⁸

And this raises a couple of apparent problems for contextualist responders to skepticism that often fly under the banner of the 'factivity problem'—problems of a

and what I'm doing in the text is stipulating that my use of 'know' will have one of the values that the term can have in ordinary, non-stipulated use.

³⁸ (c) follows from (b), given the factivity of knowledge_o and an appropriate closure principle for knowledge_h. [1. I know_h that I know_o that I have hands (Assumption for reductio). 2. I know_o that I have hands entails that I have hands (by the factivity of knowledge_o). So, 3. I know_h that I have hands (by 1, 2, closure of knowledge_h). But, 4. I don't know_h that I have hands (b). So, c. I don't know_h that I know_o that I have hands (by 1-4, reductio).] Of course, the formulation of an appropriate closure principle is a messy deal, so there may be some question about the key inference to 3, above. So let me just say here that I accept my commitment to (c), despite this messiness. (In fact, I also accept a commitment to this: Even those who don't know by any standards that knowing_o that p entails the truth of p can't know_h that they know_o facts like that they have hands. My discussion of applying skeptical arguments to the non-philosophical in Appendix D should give an indication of how I think this all plays out.)

type that, to the best of my knowledge, were first raised in print in (Williamson 2001), but have since received a lot of discussion.³⁹

Some (not Williamson) who press ‘factivity problems’ present

d. I know_h that I know_o that I have hands,

or the likes of (d) (they may substitute for the ‘I have hands’ part), which, of course, directly contradicts (c), as part of the contextualist’s position. I don’t think we contextualists have ever actually claimed to know_h that we know_o the very things we admit that we don’t know_h. However, (a) is a central claim of ours, and that apparently can lead some to take it as part of our position that we know_h that (a) is true (see Brendel 2005: 247). Another factor that might be involved in leading some critics to ascribe (d) to us is that we contextualist responders to skepticism will seem to assert the likes of (a) in contexts where skeptical scenarios have been made salient, and so seem to be contexts that we will take to be governed by standards h. For reasons we will soon see, those assertions of ours can seem to be acts by which we contextualists represent ourselves as knowing_h that (a) is so. Sometimes those who present ‘factivity problems’ just take it for granted that contextualist responders accept the likes of (d), or would at least claim that we *can* know_h that we know_o the likes of that we have hands.⁴⁰ Thus, when they (rightly) derive from us a

³⁹ See, e.g., (Williamson 2001), (Williamson 2005), (Brendel 2005), (Wright 2005), (Baumann 2008), (Brendel 2009), (Baumann 2010), (Jäger 2012), (Dinges 2014). The problem Wright presses does not presume that standards h govern the contextualist’s claim. Since the considerations of Appendix C are needed to respond, I discuss and respond to Wright in note 1 of that Appendix.

My (fallible) recollection is that I first encountered what later came to be called the ‘factivity problem’ long before this literature emerged, in work by Bredo Johnsen, as early versions of what later became (Johnsen 2001) contained this problem. In his acknowledgements, Johnsen thanks me for comments on ancestors of his paper. But one of my great personal-professional regrets in life is my fear (though I could be misremembering things) that I gave Bredo what amounted to very bad advice with respect to the factivity problem. I thought of that alleged ‘problem’ as just some kind of mistake, based on a fairly simple misunderstanding of the contextualist position on skepticism. Truth be told, I still tend to think that way, and have a hard time seeing what all the fuss is about here. But it is a fuss that has been prominently raised, and Johnsen saw the problem (or at least the appearance of a problem) long before it became a big thing. My worry is that my reaction to that aspect of his work played a role in leading Johnsen to abandon that criticism, which, in retrospect, he would have been wise to raise, and raise back in the 1990s (breaking that part of his work off as a separate paper).

⁴⁰ See, e.g., (Baumann 2008: 583), which lists at p. 581 other presentations that press a similar form of the problem.

commitment to the likes of (c), they take themselves to have uncovered a conflict in our position.

However, (d) is no part of my position—and I suspect the same is true of other contextualist responders to skepticism. I deny it. My position is that we not only fail to know_h that we have hands, but also fail to know_h that we know_o that we have hands. (c) is not some embarrassing commitment that snuck up on me, but is something I've long accepted—and have even used to my own advantage.⁴¹ It's not so much an *admission* I make, but something I *insist* on.

But denying (d) is one thing; having the right to deny it is quite another. I might be committed to (d), whether I like this commitment or not. Instead of the factivity problem being one where (d) is taken to be a part of our position that we'd happily accept, but that turns out to be inconsistent with something we contextualists turn out to be unhappily committed to (c), the problem can instead be presented as one of our turning out to have two commitments—to (c) and to (d)—that are in conflict with one another. I accept my commitment to (c), and so my hopes of escaping the problem are pinned on avoiding (d). Or, more generally, the factivity problem can be viewed as a difficulty we contextualist responders to skepticism have in stating our position, and particularly in stating (a), in the very philosophical settings in which we're likely to do so.

In either case, the potential problem can be generated by, or at least sharpened by, the Knowledge Account of Assertion (KAA), which Williamson and I both accept. KAA can be packaged as a principle concerning proper assertion, or (as Williamson likes to deal with it) as a rule governing the practice of assertion, or (as I often like to use it) as a principle concerning what a speaker represents as being the case in asserting something:

KAA-PA: A speaker is well-enough positioned with respect to p to be able to properly assert that p if and only if she knows that p according to the standards for knowledge that are in place as she makes her assertion.

⁴¹ See (DeRose 2009: 207) for how our failure to know_h that we know_o can be used by the contextualists to escape an advanced 'now you know it, now you don't' problem (the 'fortified first-person dialogue'). As I report at (DeRose 2009: 209, n. 18), I utilized this escape in talks I gave in the years 1996-2001, but since there were related 'now you know it, now you don't' problems that escaped this escape, and had to be handled a different way, this didn't come out in print until (DeRose 2009).

KAA-Rule: One must assert that p only if one knows that p according to the standards for knowledge that are in place as one makes one's assertion.

KAA-Rep: In asserting that p , a speaker represents herself as knowing that p according to the standards for knowledge that are in place as she makes her assertion.

I have presented these versions of KAA in forms which are relativized to the standards for knowledge in place as a speaker makes her assertion, a complexity that will be unnecessary to invariantists like Williamson, according to whom a single epistemic standard governs all contexts. Relativizing KAA in this way renders it neutral with respect to the debate between contextualism and invariantism: The invariantist will just suppose that the standards in place are always the same ones.

The potential problem now is easy to see. Given KAA, asserting (a) in a context governed by h will (to take the Representation form of KAA) involve me in representing it as being the case that I know _{h} that I know _{o} that I have hands—i.e., will involve me in representing (d) as being the case. But then, given (c), I'll be representing something false as being the case—and (to take the other forms of KAA) I'll be breaking the rule for assertion and making an improper assertion. All of these render it conversationally naughty to assert (a) in a context governed by h , and can be seen as a way that I generate some kind of commitment to (d) by asserting (a) in the very philosophical contexts in which I am likely to present my contextualist solution to skepticism, even if I don't explicitly assert (d).

Truth be told, I've never seen much of a problem here. The lesson for contextualist responders to skepticism would just seem to be: If you want to flat-out assert (a) (without being naughty and without generating a conflict with other essential elements of the contextualist solution to skepticism), you should avoid allowing h to come to govern the contexts in which you are speaking. But that's not really a special problem for asserting (a): Given KAA, if one wants to non-naughtily flat-out assert *much of anything* (including, notoriously, any controversial skeptical claims), one should avoid allowing h to come into force, since we know little-to-nothing by standard h .

And such avoidance is what I, for one, at least try to achieve. I do not acquiesce in the skeptic's 'attempts' to raise the standards to h .⁴² When the skeptic

⁴² 'Attempts' is in scare quotes because the skeptic may not construe herself as trying to do any such thing: She may think h always governs every conversation, and that she is just trying to get me to realize that. But

presents AI, I don't 'accommodate' her and start speaking as if I know little-to-nothing. I instead engage in verbal behavior which constitutes resisting her 'efforts', and at least trying to 'not let her get away with' raising the epistemic standards to h.⁴³

Complications:⁴⁴ Now, there could be questions about whether I can succeed in my attempts to keep h from being installed. That depends on just what the rules are for which epistemic standards govern a context when the speakers involved are 'trying' to push those standards in different directions. I canvassed several contextualist options on this matter in SSS (see note 11, above). Some presentations of the factivity problem clearly assume very 'skeptic-friendly' versions of contextualism (by which I mean versions on which the raising of epistemic standards happens easily, often even when that means overriding conversational maneuvers that would tend to keep the standards from so rising), though it's sometimes difficult to tell.

But let's first see how this plays out on Williamson's presentation of the problem, which, interestingly, avoids reliance on KAA by instead (though this is closely related) charging that the contextualist responder to skepticism gets committed to a Moore-paradoxical statement:

Once sceptics have manipulated the context, in the epistemology seminar, contextualists are apt to console themselves with the thought that although most denials of "knowledge" in that context of scepticism are correct, in everyday contexts many assertions of "knowledge" are also correct. For example, although 'I do not know that there is a whiteboard in the room' expresses a truth as uttered in the seminar, 'He knows that there is a whiteboard in the room' expresses a truth as uttered simultaneously outside the seminar about the same person. But that thought underestimates the gravity of the situation in which the sceptic has put contextualists, on their own analysis. For since 'know' is a factive verb, the truth of 'He

however she views what she is doing, by contextualist analysis, she is executing maneuvers which have a tendency to raise the epistemic standards, and so she can be viewed, in an extended sense, as 'trying' or 'attempting' to raise the standards. See the related discussion above in section 5.

⁴³ I myself tend not to do so by insisting that 'I know that I have hands', or the like, but what I do does clearly constitute resistance to the skeptic's attempted raising of the standards.

⁴⁴ The above material constitutes the essence of my response to the 'factivity problem'; many readers who are not particularly obsessed with that problem may wish to skip the complications below.

knows that there is a whiteboard in the room', as uttered outside the seminar, requires the truth of 'There is a whiteboard in the room' as uttered outside the seminar and therefore as uttered inside the seminar too (since no context-dependence in the latter sentence is relevant here). Thus the consoling thought commits them to claims such as 'There is a whiteboard in the room', while their contextualism commits them to claims such as 'I do not know that there is a whiteboard in the room'. Consequently, in the seminar, they are committed to this:

(MK) There is a whiteboard in the room and I do not know that there is a whiteboard in the room.

But (MK) is Moore-paradoxical: although (MK) could be true, it is somehow self-defeating to assert (MK). Thus contextualists are not entitled to the consoling thought. In the seminar, they should not say that 'He knows that there is a whiteboard in the room' expresses a truth outside the seminar At least in the seminar, contextualists are at the sceptic's mercy and cannot look outside for help. (Williamson 2005: 689)

Williamson's assessments that 'contextualists are at the sceptic's mercy' and that our contextualism 'commits [us] to claims such as' the second, negative half of his (MK) may well be based on the assumption of a skeptic-friendly version of contextualism, that we have seen is optional, and is not in fact my view. But it's not easy to interpret the claim that I am 'committed' to the likes of 'I don't know that I have hands'. Perhaps the idea is that the rest of the above passage is supposed to be understood under the assumption, set at the start of the quoted material, that the skeptic *has succeeded* in installing *h* as the standard that will govern *my* claim, in which case, if I am going to say either 'I know that I have hands' or 'I don't know that I have hands', then it has to be the latter if I am to avoid saying something false. But if I, as the other party to the conversation, have any say in whether the skeptic succeeds in installing *h*, as I do on many not-so-skeptic-friendly versions of contextualism, we can't just assume that the skeptic succeeds. So perhaps the lesson we should take from all this is that, insofar as we sense that it's permissible to stick to one's guns against a skeptic, we just shouldn't opt for such a skeptic-friendly version of contextualism—and then also that we should resist skeptics. This would connect with the main lesson to be drawn of the 'factivity problem' if we do endorse KAA—namely, the wise piece of strategic advice that, assuming KAA, it is foolish to complacently let the skeptic drive the standards up to *h* on the grounds that, having given up hope of being able to truthfully say the simple 'I know that I have hands', you can still always fall back on being able to flat-out assert with perfect propriety the likes of (a). This is foolish, given KAA, because (a)'s commitment to the truth of 'I have hands' means that (a) too is unassertable in contexts governed by *h*.

But back to Williamson's 'commitment'-to-MK version of the problem: There are other ways to understand the charge that I am 'committed' to the likes of 'I don't know that I have hands' when talking to the skeptic. One is just a small tweak away from the interpretation we tried above. We can take it to be the claim that if I am going to say either 'I know that I have hands' or 'I don't know that I have hands', I am committed to it having to be the latter if I am going to say something true (as opposed to just avoiding saying something false). In this way, I am committed to 'I don't know that I have hands' when talking with the skeptic on a wider range of versions of contextualism—including, interestingly, the 'gap view' which I myself lean toward (DeRose 2009: 144-8), for on that view I will succeed in saying something true if I agree with the skeptic that 'I don't know that I have hands', but will say something that is neither true nor false if I insist that 'I know that I have hands' (in the face of a stubborn skeptic). We may be tempted to take this as a reason to opt for a version of contextualism still less friendly to the skeptic, like the 'veto power' view (DeRose 2009: 140-41), or perhaps a 'multiple scoreboards' view (DeRose 2009: 134-5), on which I will assert a truth if I insist that 'I know that I have hands', and so on which I am not in the way under consideration 'committed' to 'I don't know that I have hands' (and on which, in a good sense, I am not so 'at the skeptic's mercy').

But I for one don't see this as a very strong motivation for moving to such a version of contextualism, mostly for reasons we'll see a bit later, but also because our current understanding of the term is a rather strange way to be 'committed' to the sentence, 'I don't know that I have hands', for, among other problems, it seems to assume that I must assert either 'I know that I have hands' or 'I don't know that I have hands', and where is that forced choice coming from? It doesn't seem that I would have to be problematically unforthcoming in opting to assert neither of those two things. When pointedly asked whether, according to me, I know that I am not a BIV, I could for instance instead assert (a), at which point I'd feel I've pretty well described my position on the issue at question. Of course, given KAA, as we've seen, it's problematic to assert (a), too, in a context governed by h. But this may perhaps show that one really does need to rely on KAA to drive a 'factivity problem', and we should give up trying to foist some problematic 'commitment' to the sentence 'I don't know that I have hands' onto the contextualist responder to skepticism.

But perhaps the way in which I'm committed to 'I don't know that I have hands' is supposed to not be anything like that I'm in some way forced to say it, but

just that I'm committed to what is asserted being true *if I*, or if some contextualist responder, (or perhaps even if the skeptic), *were to* assert that sentence in a context in which skepticism is being pressed (since I'd then be going along with skeptic's attempt to raise the standards). And the contextualist would have that kind of commitment to the sentence on a wide variety of forms of contextualism. (Not all: some ways of working out a 'reasonableness view' (DeRose 2009: 141-2), for instance, would allow us to avoid even this kind of 'commitment', as would certain forms of contextualism which place a 'ceiling' on how high the standards can go.) But it's an odd complaint that the contextualist is 'committed' in *that* sense to the sentence 'I don't know that I have hands', and then to the Moore-paradoxical sentence, (MK), of which it is a part. For, as Williamson points out, and is clear in any case, Moore-paradoxical sentences like (MK) could be true; that's really the interesting thing about them. If you want to avoid making Moore-paradoxical pronouncements, then, contextualist responder to skepticism or not, you're going to have to refuse to say some things that would be true for you to say. So it's not much of a complaint against someone that they're 'committed' to a Moore-paradox in a sense which requires for such a 'commitment' to a sentence only that one's views imply that one would speak a truth in asserting that sentence.

And if the idea is just that I'm committed to the truth of 'I don't know that I have hands', and then to (MK), when they're evaluated at *h* (nevermind who is or could be saying any of this), then I have no problem with having such a commitment to a Moore-paradoxical statement. Some Moore-paradoxical statements are true, and I'm perfectly happy to be committed to (MK), as evaluated at *h*, being one of them. In fact, I wouldn't have it any other way.

So, let's stop worrying about 'commitments', and go back to KAA-generated factivity problems. And what we seem to be left with is a potential problem, generated by KAA (which I accept), for contextualist responders to skepticism asserting our views in contexts governed by *h*. To which, again, the solution would seem to be: Don't let your assertions be governed by *h*—which is very good advice in any case! The remaining potential concern would then seem to be the worry that the very philosophical contexts in which contextualist responders present their solutions can appear, at least on a certain view of how contextualism operates, to be ones which are or maybe even must be governed by *h*.

So, we are faced with a couple of (closely related) questions: On the best versions of contextualism, are contextualists' assertions of our view typically

governed by *h*? And: On the best versions of contextualism, can the skeptic force us into a context where our assertions are governed by *h* (are we ‘at the skeptic’s mercy’)? And let’s just suppose that it would be a good objection to contextualist responses to skepticism if the answers to these questions were yes.

For the answers to these questions are: no and no. And quite clearly so, this old contextualist would say, even though we contextualists typically find ourselves asserting elements of our responses to skepticism in contexts in which skeptical hypotheses have been mentioned and are hanging in the conversational air. But so far, I’ve been citing the various options for contextualists of the type I canvass in SSS. And, yes, it’s very fair to say that I (and other contextualists, too) haven’t been very definitive in which of those options we think is best (though I and we do have our leanings). And some of these options—including the one I lean toward, as we saw above—do have potential for the skeptic to be able to conversationally push us around in some ways.

But the choice among the options I present in SSS there occurs against the backdrop of a certain kind of dispute between skeptics and those who resist them, disputes falling in the parameters set at (DeRose 2009: 131-4). And in presenting our views, contextualists don’t have to stay within those parameters. And if you are worried that I am often in fact subject to factivity problems in asserting my views, never fear: I in fact don’t stay so boxed in. In particular, in SSS I was addressing disputes where no party to the discussion offers a ‘contextualist, or proto-contextualist, analysis of the debate that is transpiring’! And rest assured that when I am presenting my contextualist treatment of skepticism, I of course do not—at all—stay within *that* parameter. And if it comes to this, once a party to a conversation starts explicitly stating what value their use of a context-sensitive term is to have, we are of course in a whole different situation. If the other party refuses to go along with using the value specified, then we are in a conversational situation where a ‘multiple scoreboards’ treatment seems very much in order, in which case the contextualist responder to skepticism is clear of the ‘factivity problem’, for *his* assertions will not be governed by *h* (though the skeptic’s assertions, governed as they are by *h*, *are* still in grave danger).

So, yes, there can be tough and very tricky calls for the contextualist (or the non-contextualist, trying to judge which are the most viable forms of contextualism) to make concerning the semantics of disputes taking place ‘in the dark’ about the possibility of contextualism, where the parties to the dispute, though so in the dark,

are still ‘trying’ (in the sense we have been using) to push the epistemic standards in different directions. *Maybe* (though this goes against my own inclinations) a very skeptic-friendly treatment of such situations is in order. If so, then in a good sense the unenlightened Moorean is at the skeptic’s mercy—and, given KAA, at her mercy with respect to her flat-out assertions generally, and not just with respect to her explicit claims to knowledge. But that wouldn’t mean that (on contextualism, or on the most viable versions of it) *the contextualist* who realizes what is going on is at the skeptic’s mercy. If it really comes to this (and I don’t think this is really necessary), and you don’t want your claims to be governed by *h*, just say so.⁴⁵ That really should take care of matters.

A final note: Central to this section has been my embrace of

c. I don’t know_h that I know_o that I have hands.

That embrace is sincere, but I worry that it, and my attention to the conditions under which standards *h* would govern a discussion of skepticism, may give the impression that I think that so long as I keep the standards for knowledge from spinning out of control, I’m in the clear to assert my philosophical views, despite my acceptance of KAA, because I do take myself to know these views to be true by ordinary or moderate standards for knowledge. But that is not so. I don’t take myself to know even by ordinary standards that my contextualist solution to skepticism is correct.⁴⁶ For discussion of this ignorance_o (lack of knowledge-by-

⁴⁵ As a look at SSP will reveal, I don’t there include any explicit instructions that I mean not to be governed by standards like *h*. That’s because I don’t think such pedantry is needed there. Part of the reason I suppose is that I don’t construe my reader as a skeptic, but just as someone who is interested in the issue of skepticism. Thus, though I am up to my eyeballs in mentions of radical skeptical hypotheses, I don’t take my own claims to be there governed by *h*. That I am engaging in contextualist analysis, and making it clear that I don’t take myself to know_h much if anything, while I do take myself to know much by other standards, suffice to make it clear that my claims are not to be understood as governed by *h*, I hope. When the situation does call for it (as when I am facing a real skeptic who is pushing the standards toward *h*, and especially if she (or if some non-skeptic) starts accusing me of having ‘factivity problems’, or if I’m asked whether I take my claims to be governed by *h*), I will say that I don’t take my talk to be governed by *h*. And if it helps, I explicitly declare that right now. Those seem very bad standards to use!

⁴⁶ That contextualism is true, and that my contextualist solution to AI skepticism is (even roughly) correct are both the kind of controversial philosophical claims that I don’t take myself to know, even by ordinary standards. But that I know_o that I have hands (which does not depend on contextualism, since it can be true even if invariantism is correct) may be the kind thing which we can, and I do, know_o, even if we can’t know_h it.

ordinary-standards), including whether it gives rise to something in the vicinity of a 'factivity problem',⁴⁷ see Appendix C.

My worry here is that I'll give readers to think that I take myself to know, even my controversial philosophical claims to be right.

⁴⁷ See n. 1 of Appendix C.