

Appendix C: Do I Even Know_o Any of This to Be True?: Some Thoughts about Belief, Knowledge, and Assertion in Philosophical Settings and Other Knowledge Deserts

At the end of Chapter 4, I addressed the worry that I face some special ‘factivity problem’ in asserting the philosophical views I’m defending in this book. The alleged problem was largely driven by my admission/commitment to my lack of knowledge_h of a key component of my view, and I spent some space discussing the conditions under which standards_h would govern a discussion of skepticism. I fear that my discussion may have given 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.

And honesty compels me to say that’s just not so. I don’t take myself to know even by ordinary standards that my contextualist solution to skepticism is right. And while this generates issues concerning how I am in a position to state my views, these turn out to be very general issues in stating philosophical views—general issues that I will quickly address here.¹

¹ This is important to acknowledge not only to avoid giving a false impression, but also to answer one form of what may be classified as a ‘factivity problem.’ In presenting what is classified as a version of that problem, Crispin Wright distinguishes three contexts: ‘Consider three epistemological contexts: that of quotidian common sense, Q; that of scepticism, S; and that of epistemological contextualism itself, C’ (Wright 2005: 243). In pressing his problem with the contextualist solution, Wright does not assume any position on whether C is distinct from S. Instead, Wright alleges this: However C aligns (or fails to align) with the other contexts in question, if the contextualist admits that she does not know that her position is correct by the standards that govern C (the context in which she presents her view), she is thereby committed to having no point to make on these matters in C: ‘Now the perspective of epistemological contextualism had presumably better be one from which the theorist can know both of the potentialities of quotidian common sense and of the limitations imposed by skepticism—can know both that common sense’s claims to knowledge are, in their proper context, perfectly good, and that scepticism’s denials of knowledge are perfectly good in theirs. For if contextualism cannot rationally profess that knowledge, it has no point to make!’ (Wright 2005: 243). I am here admitting that I don’t even know by ordinary standards (I don’t know_o) that my contextualist solution is right, and so I’ll also admit to Wright that I don’t ‘know_c’ that my solution is right—I don’t know that it’s right by the standards for knowledge that govern the contexts in which I present it. But as will emerge in this

The first thing to say here is that I do not at all feel alone in my predicament: I think philosophers generally don't know—by any good standards—the positions we take on controversial issues.² (And, yes, that too is something I don't take myself to know, by any good standards.) If that renders us pointless (see note 1), we are in trouble! In fact, I think there is something to the view of philosophers as specialists in addressing some of the questions we find important, but which nobody has yet figured out a knowledge-producing way to get answers to, generating answers to such questions, and good (even if not knowledge-producing) support for those answers. If that is a big part of what we are good for, then 'Stick to the points you know to be right' would be about the worst advice any philosopher could follow in their work!

Though I hold that we philosophers typically don't know that our controversial positions in philosophy are correct, even by the ordinary standards for knowledge at which we count as knowing lots of other things (and for the rest of this Appendix, my uses of 'know' and its cognates should all be understood as designating such knowledge.), it can often *feel* to us as if we do know that we are right. And this can give rise to delusions of knowledge.

Truth be told, I think that we don't even really *believe* that our controversial views in philosophy are correct—though I should quickly clarify that, and particularly the 'really' I throw in there. While I feel quite comfortable in judging that we do not know the items in question, when it comes to belief, I think that we are typically in an 'in-between' state, as Eric Schwitzgebel puts it in his very helpful work here (Schwitzgebel 2001; 2002), in which it wouldn't be true to describe us either as 'believing' or as 'not believing' these things. My use of (the quite elastic word) 'really' here is intended so that 'really believing' would be to be in a state in

Appendix, I think philosophers *generally* don't know that their positions are correct (by the standards that govern the contexts in which they present those positions), and I insist that we—and I—do have points to make nonetheless.

² As Gary Gutting points out, while philosophical theorizing does not typically result in knowledge of the controversial positions we take on the issues we're focused on, it does produce, as a regular byproduct, philosophical knowledge of such things as that there are certain important distinctions to be drawn (Gutting 2009). I'm not focused here on getting the line between these known byproducts of philosophical theorizing and our philosophical positions, which I'm claiming we don't know, exactly right. I would be in rough agreement with Gutting about how to do that.

which one could be truthfully said to believe the item in question.³ In this sense, we do really believe all sorts of things, including, I think (since I'm generally quite generous in ascribing beliefs in lots of other, non-philosophical cases), things we are quite unsure of. Example: I hold that in the intended sense I do really believe that Abraham Lincoln was born in 1809, though I'm unsure of that. I seem to remember that being the year given by Goldman in an example in a paper (Goldman 1967) that I have in the past read and taught from. However, I'm now uncertain that 1809 was the year used in the example, and so whether it is really the year Lincoln was born in. If someone were to ask me what year Lincoln was born in, I wouldn't quite feel in a position to flat-out assert that it was 1809, but would only give a hedged answer, e.g., 'I *think* it was 1809.' Yet, despite taking us as really believing even things that we are unsure of, I *still* think (and here, 'think' conveys philosophical acceptance: this too is something I don't really believe), we don't really believe our controversial views in philosophy—though neither do we really not believe them.

Such thoughts seem fairly common among philosophers nowadays, but I will illustrate the kind of considerations that lead *me* toward them by using the example of the philosophical issue that William Alston used to make somewhat related points (Alston 1996: 10-11): So, I'm an incompatibilist about free will and determinism. This is a view I accept and will defend, sometimes passionately, in various settings of philosophical discussion. This is a good example for me to use here because, not only is this outside of my areas of main expertise, but it is a view on which the majority of philosophers, and also the majority of philosophers who have studied the issue much more closely than I have, seem to be lined up against me.⁴ *Still*, in

³ This particular use of 'really' is useful to those who do not accept bivalence (about whatever claims are in question). Not being at all a fan of bivalence, I find it quite helpful. When it comes to belief, I don't know which side I'd come down on if I really (different use!) had to choose between 'We believe' and 'We don't believe' (our controversial philosophical views that we feel very strong about) here. This is one place where I'm especially happy to have the middle ground available—even though that means I have two boundaries (between 'yes' and 'kinda' and between 'no' and 'kinda'), rather than just one, to worry about.

⁴ Bryan Frances has done important work in getting philosophers to think about the challenge of facing disagreement by our 'epistemic superiors', as well as by our 'epistemic peers'. See, e.g., (Frances 2012).

In the 2009 PhilPapers Surveys [<http://philpapers.org/surveys/>], the results for 'Free will: compatibilism, libertarianism, or no free will?' were:

Among 'All respondents' in the 'Target faculty':

Accept or lean toward: compatibilism 550 / 931 (59.1%)

the philosophical settings in which I sometimes find myself contending for incompatibilism, I nonetheless *feel* strangely confident that I'm right. It indeed *feels* to me very much like something I know to be the case—and *certainly* like something I believe to be the case.

But things would be very different if actual practical consequences were somehow tied to getting this matter right. Suppose I'm up on the ship of super-advanced aliens, whom I somehow (and nevermind how) know to be truthful when they tell me that the issue of whether free action is compatible with determinism is one of those philosophical questions we humans puzzle over that actually does have a correct answer; that they, the aliens, actually know whether it's compatibilism or incompatibilism that is correct; and that they will give me a chance to save the Earth and humankind by getting the question right: I get to give one answer to the question whether compatibilism or incompatibilism is true, and if I refuse to answer or get it wrong, they will destroy the Earth and everyone living there, but if I get it right, they will destroy nothing, but will return me to Earth and then peacefully leave. Or, to vary the case in a couple of different ways, suppose first that it is not the fate of the Earth and humankind that is at stake, but only my own life⁵; or, secondly, that no lives are at stake and my entire encounter with the aliens is a very friendly

Other	139 / 931 (14.9%)
Accept or lean toward: libertarianism	128 / 931 (13.7%)
Accept or lean toward: no free will	114 / 931 (12.2%)

Among specialists in 'Metaphysics' in the 'Target faculty':

Accept or lean toward: compatibilism	130 / 234 (55.6%)
Accept or lean toward: libertarianism	50 / 234 (21.4%)
Accept or lean toward: no free will	31 / 234 (13.2%)
Other	23 / 234 (9.8%)

Among specialists in 'Philosophy of Action' in the 'Target faculty':

Accept or lean toward: compatibilism	23 / 43 (53.5%)
Accept or lean toward: libertarianism	8 / 43 (18.6%)
Other	7 / 43 (16.3%)
Accept or lean toward: no free will	5 / 43 (11.6%)

⁵ The save-the-Earth case raises the worry (or I suppose a pack of related worries) to the rough effect that I defer to the experts here in order to avoid, at least to some extent, taking on myself the responsibility for the lives of others. Some of these worries are met by varying the case so that it is only my own life that is in jeopardy, which keeps the stakes quite high, while avoiding doing so by making me responsible for the lives of others. Thanks to Zoe Kempf-Harris for suggesting this case variant.

and positive one, but that the aliens will instead give me 10 million U.S. dollars if I give the correct answer, but nothing if I'm wrong, before releasing me and peacefully leaving. In any of these cases, I'm sure I would feel very differently about the issue than I do when discussing the matter in a philosophical setting. And what's really interesting is that, beyond the effects one would likely expect high stakes to have on the matter, at least until recently, I would have been strongly inclined to go with the opinion of the majority of philosophers, rather than my own philosophical acceptance of the matter, in these cases.⁶

I realize that it's quite dicey to predict what one would do in such wild circumstances, and I have now encountered (when presenting these thoughts at various places) lots of interesting guesses others have made about what they would do in the relevant situations (I asked listeners to adjust the philosophical question to one about which they feel strongly but find themselves at odds with the majority of expert opinion), and interesting opinions about what it would be rational to do. But so long as I would feel at least a significant temptation to 'flip' (to go with majority expert opinion, rather than with how things seem to me personally), this seems to be in marked contrast to how real beliefs, even those held fairly tentatively, behave under such stress. Consider again my uncertain belief that Lincoln was born in 1809. Since I'm unsure about the matter, I will be quite conservative in what I'll stake on that belief. But in situations in which it's clear that I should give an answer to the question (like when something bad will happen if I refuse to answer or give a wrong answer, and will be avoided only if I answer correctly; or, positively, if something good will happen only if I answer and answer correctly), 1809 is the answer I'm giving, with no temptation to opt for a different one. If you greatly raise the (positive or negative) stakes on me, you can make me feel in various ways very unsure about what will happen. When there's a huge negative result on the table, you can cause me great anxiety. But you won't thereby tempt me to go with 1808 or 1810 instead of 1809. You may make me add the likes of 'Heaven help me!' to my answer. And you may make me try very hard to search

⁶ This example has been somewhat ruined for me since I have dipped a bit more deeply into the philosophical literature on free action, and have become more and more convinced of how little the compatibilist position has going for it. This complicates matters greatly, as it is now quite impossible for me to confidently guess what I would do in the save-the-Earth and in the save-my-own life cases, and makes me now think I'd go with my own incompatibilism when it's a million dollars that's at stake. But rather than deal with these complications, let's just focus on my earlier self, who was not so well apprised of the state of the discussion.

my memory more carefully. But insofar as all I can come up with is this push toward the answer '1809', that's what I'm going with, and that is in its way an easy (even if anxiety-producing) call for me to make: while very worried, I am not at all tempted to flip.

And I remain untempted to flip even if I have some ordinarily dominated push toward giving another answer: that push that is dominated in ordinary circumstances remains dominated as the aliens dial up the pressure.⁷ Suppose that in addition to my not-so-definite recollection of 1809 being the year used in the Goldman paper, I feel some push to go with an earlier year to make better sense of how old Lincoln looks in a particular picture of him that I know was taken in 1863. (If it helps, you can suppose that I have a choice between 1809 and, say, 1806, and I feel that 1806 makes better sense of how old Lincoln looks in the 1863 picture.) Adding that conflict increases my anxiety in the high-stakes situation, but doesn't tempt me to change my answer from the one I would give in a low-stakes situation. I'll perhaps consider the matter more carefully, but if all I come up with is the same two pushes, I'll almost certainly judge the relative strength of those two pushes the same as I do in the low stakes scenario and go with 1809.

But things are very different in the case of my incompatibilism (which, at least in some heated philosophical settings, doesn't feel very tentative at all). There too, there are indications pointing in different directions: how the issue and arguments strike me personally points toward incompatibilism; the weight of expert opinion

⁷ Since several very smart readers/listeners have felt moved to point out to me here that I would probably be tempted to flip on the year of Lincoln's birth if I ran up against credible enough disagreement on that while aboard the aliens' ship, it's worth saying that I very much agree with that point. What I take to be the important difference between my attitude toward the year of Lincoln's birth on the one hand and toward my controversial philosophical views on the other is not how they 'behave', generally or in particular in response to disagreement, when something is actually riding on my being right about them. What's weird (at least relative to the assumption that they're real beliefs) is how my philosophical positions behave in the seminar room, where they seem oddly unmoved by peer disagreement, or even disagreement with those who seem better positioned than me. This oddness gets exposed by what happens on the aliens' ship, as I'm tempted to flip away from the weird attitude I display in philosophical settings (in the seminar room). As I here consider what happens when I add some dominated push against my belief in the Lincoln case, it's important that it be a *dominated* push, and the push that credible enough disagreement would provide would not be dominated—on the alien's ship, or in everyday life. I suppose not-so-credible disagreement (e.g., issuing from someone who seems even much more uncertain than me) could work here: the weaker push provided by such disagreement would be dominated in both everyday life and under great practical stress (though it would of course occasion more feelings of anxiety in the high-stakes situation).

points toward compatibilism. But on that issue, raising the stakes does have a marked effect on the relative weight I assign to those two indications: expert opinion, which has little-to-no effect on me when arguing about the issue in the seminar room suddenly becomes a very weighty consideration up on the aliens' ship. Consequently, I will be very tempted to flip my answer, and give an answer on the ship different from what I give in philosophical discussion.

Do such facts about how I would have acted when something was actually riding on whether I was right about the matter mean that (until recently) I really believed that compatibilism is true? I think not, but I do think that this and related considerations do point to the conclusion that I didn't really believe that incompatibilism is true—despite how strongly I might have felt that I believed it when arguing about it in philosophical settings.⁸

The notion of believing something to be the case seems to be somehow closely tied to various dispositions: dispositions to act in certain ways under certain circumstances, to say certain things under certain circumstances, to have certain emotions under certain circumstances, to form certain other beliefs under certain circumstances, etc. And my suspicion that we don't really believe (and also

⁸ Compare William Lycan, whom I've recently discovered has very similar thoughts about belief:

First, there is a question of what philosophers believe. Do we believe our own views in the same sense in which we believe ordinary things? Hume and Moore pointed out that there is often a very striking gap between our doctrines and our behavior: in everyday life, epistemological skeptics make knowledge claims, idealists kick stones, antirealists about time consult their watches, and so on. (A similar point has often been made about religious believers, many of whom do not behave at all as if the creeds they profess are true.) It is important to see that these are typically not cases of hypocrisy or other insincerity; when we avow a philosophical thesis and defend it, sometimes passionately, we mean what we say. Nonetheless the thesis does not show in our nonverbal behavior, or even in our off-duty verbal behavior.

Also, suppose I have been thus passionately defending one of my own core philosophical doctrines, say functionalism regarding the mental. I have said in print that functionalism is a view that I would kill or die for. But now suppose I encounter an oracle who knows the truth of the matter, or perhaps God parts the clouds and tells me that in sixty seconds He will reveal whether or not functionalism is true. You invite me to bet \$1,000 on functionalism. Would I take that bet? Of course not. Nor \$10 for that matter. Are you nuts? This is philosophy we are talking about. In reality, I have no idea whether functionalism is true. (Lycan 2013: 115)

Gutting, with whom Lycan is interacting here, seems inclined to agree with Lycan about this lack of belief (Gutting 2013: 135).

that we don't really not believe) our philosophical positions is best understood in terms of the dispositions involved in believing: We lack key dispositions that would be tied to our being real believers, but we also have dispositions that make it untrue to deny that we believe our positions.⁹

It may seem that I am putting too much weight here on our lacking dispositions to act as if our philosophical positions are true in certain very unlikely scenarios. But far-fetched as my advanced aliens scenario is, how we would respond to situations like that seems nonetheless quite important to the issue of whether we really believe those positions. The reason for my focus on such an unlikely scenario is that it is one in which something actually turns on whether we are right. In realistic situations, we are not subject to actual consequences to our being right or wrong about our philosophical acceptances. And what kind of real beliefs are these if we would 'flip' on them as soon as something actually turned on our being right?¹⁰

As I've stressed, one of the sharpest contrasts in my thinking about the issues in the imagined far-fetched, high-stakes situations as opposed to the philosophical settings in which I actually think and talk about the issue is the role of expert philosophical opinion. It would weigh very heavily in my deliberations in the imagined high stakes scenarios. And importantly, its effect would not be limited to making me more hesitant about my answer, but it would have the power to actually flip what answer I give. In my actual thinking in philosophical settings, by very sharp contrast, and quite suspiciously, the expert opinion of other philosophers carries at most hardly any weight at all.

And though that imperviousness to disagreement fuels my suspicion that these are not real beliefs, it is for the good, I think. It's probably good for

⁹ For helpful related discussion, see esp. Schwitzgebel 2002.

¹⁰ I chose an example in which I would actually 'flip' in the high-stakes scenario. But my suspicion also covers cases in which, for instance, one's own philosophical acceptance is aligned with the majority of expert opinion, so one would not flip, but would nervously stick with one's own inclination when backed into a corner by the aliens. Such cases are better candidates for being real beliefs, I admit. But of course, there are other, easily constructed, modified alien-like situations in which one would fail to act on such an acceptance—in ways that I think cast doubt on whether it is a real belief. To seriously assess whether such acceptances are real beliefs would probably require dealing with the issue of what level of confidence a subject must have for an ascription of belief to that subject to be true—which I suspect is an elusively context-sensitive matter.

philosophical progress (which I do believe in, despite my conviction that it typically does not lead to knowledgeable answers to the questions we are focused on) that in the settings in which we think about and discuss philosophy we do not let considerations like contrary opinions by peers and experts make us go all wishy-washy about philosophical issues, or worse still, quickly flip from what seems right to us personally to what seems to be the majority of expert opinion, but instead, in a way, sincerely feel confident about the positions that seem right to us personally, passionately defend them, etc.¹¹

Let me briefly interject how I think this should effect the recently hot discussion of the epistemology of disagreement. Many of the cases used in the literature where it intuitively seems that one should, or that it's at least permissible to, 'stick to one's guns', as opposed to 'conciliate' when one encounters 'peer disagreement' over an issue concern areas of controversy, like philosophy. I suspect that the reason it's alright to stick to one's guns in such examples is that these aren't real beliefs to begin with. (And, relatedly, I think the credence we already really assign to the propositions in question is not nearly as high as one would be led to think by how confident we seem to be, and in ways feel ourselves to be, about those items when we're engaged in philosophical discussion.) It's alright to stick to one's guns in these cases only because they are only toy guns, as it were.¹²

¹¹ (Barnett ms.), which I became aware of after writing this Appendix, takes a similar view to that of this paragraph. Barnett is not focused on descriptive matters of what attitude we do take toward our contested philosophical positions, but on an evaluative question, arguing that we are not *rational* in holding such beliefs—and I certainly agree that such beliefs would not be rational. And where it comes to the value to philosophy of our accepting what seems right to each of us personally, without due regard to what others think, our views are very similar. Barnett advocates that we not believe our contested philosophical positions, and instead adopt an attitude of 'disagreement-insulated inclination', accepting those positions that are supported by one's evidence *other than the evidence involving agreement and disagreement with others*. I think that that's pretty much what we do already.

(Fleisher ms.), which I became aware of even more recently, also stresses the importance of our forthrightly contending for our views in areas of cutting-edge research (including philosophy) in the face of disagreement and our limited epistemic standing with respect to our positions—which limitation Fleisher agrees would make belief irrational. Fleisher proposes that we take an attitude that he develops and calls 'endorsement' toward our philosophical views, and holds that rational endorsement is and should be the norm for assertion in research contexts.

¹² In the (more jarring) cases of peer disagreement where real guns are in play, I think the rational response is to revise one's credence in the underlying matter and/or (though most cases will involve at least some of

Turning now to the matters of knowledge and assertion: In these settings in which we are passionately defending things we are not even close to knowing to be true, we find ourselves, among other things, flat-out asserting such things, in flagrant violation of the knowledge norm for assertion. This is also probably to the good. In philosophy and other ‘knowledge deserts’, as we might call them, where we’re focused on questions that none of us knows the answers to, it would be quite a drag to have to be constantly hedging our assertions. So, often enough, it seems, we don’t. There are great differences in personal style among philosophers, some humbly hedging their claims, by, e.g., throwing in parentheticals of the likes of ‘I think’, where others more boldly assert away, hedges be damned. But most will find themselves asserting from the hip in at least some philosophical settings. (In my observation, this most often happens in at least moderately heated philosophical disputes, where some kind of bilateral escalation in projected confidence often seems to occur.) And when it does happen, the resulting assertions don’t seem wrong—or at least, not in the way that one is wrong to assert ‘There’s a service station round that corner’ to a motorist when one is nowhere near to knowing that to be so. In philosophical discussion, we seem to have some kind of license to assert without hedges the things we accept, even though we don’t know them to be so, and to, at least in that way, act as if we know things we don’t really know.¹³

I’m very open to different ways of understanding how this license works. Indeed, some may wonder whether the speech acts in question really ought to be classified as unhedged assertions, as I have done above. In their (2014), which is a defense of contextualism from the factivity problem, but then, as I am doing here,

both) to revise one’s estimate of the reliability of one’s interlocutor, in at least roughly (and maybe more exactly) the way proposed in (Worsnip 2014).

¹³ For an alternative account of assertions in philosophy and other areas of controversy, see (Goldberg 2015, Chapters 9 and 10). Goldberg handles these assertions by means of a particular way of making the epistemic norm for assertion context-sensitive. I am of course extremely open to such context-sensitivity, as it plays a crucial role in my version of the Knowledge Account of Assertion (see esp. DeRose 2002: 181-2 ≅ DeRose 2009: 98-9), but have long thought that we cannot account for the assertions we are here addressing just by supposing the contextually-relevant standards for assertion are low in these cases, but must resort to something like my current proposal that we engage in a pretense of meeting standards that we don’t in fact meet. There are no doubt many and complex causes for this difference between Sanford and me, but I suspect that a key root of our difference is that, while we both recognize that our epistemic position with respect to our controversial philosophical views (and our opinions in other knowledge deserts) is quite limited, I take a significantly even dimmer view of our epistemic position with respect to such matters than Sanford does.

broadens out to consider what philosophers are generally doing when they put forward their controversial views, Martin Montminy and Wes Skolits (henceforth 'M&S') construe the claims in question as something weaker than assertions:

A key assumption of the argument generating the stability problem is that while in [the context in which she presents her theory], the contextualist *asserts* the content of her theory. But one may plausibly hold that the contextualist's utterances have a slightly weaker assertoric force than assertions do. Consider the category of illocutionary acts called *weak assertives*, which includes conjectures, guesses and hypotheses. These illocutionary acts aim at truth, but their assertoric force is weaker than that of an assertion. On the current proposal, the force of the contextualist's weak assertives would be somewhere in between the force of a conjecture and that of an assertion. Their illocutionary force would be comparable in strength to that of the weak assertives generated by a parenthetical use of 'I think'. . . . This strikes us as a plausible description of what typical philosophers do when they defend their views, except that they tend to avoid stylistically frowned upon parentheticals. In a philosophical context, it is understood that many of the claims made are highly controversial and cannot be established decisively. There is thus an implicit understanding that speakers do not represent themselves as knowing the content of every utterance they make. Utterances expressing controversial philosophical views are thus reasonably interpreted as having weaker assertoric forces than assertions do.

Weak assertives are governed not by a knowledge rule, but by a weaker epistemic requirement: a weak assertive that p is epistemically appropriate only if the speaker has *some* evidence for p. (M&S 2014: 327)

M&S are clearly being driven by the same kinds of considerations I am, and are seeking to end up at roughly the same destination: some kind of license in the relevant settings for the unhedged speech acts in question. However, I am not myself inclined to get to this destination by denying that the claims in question are assertions. This is mainly due to my having a somewhat different view than M&S do on some of the relevant conversational phenomena. Most notably, M&S write:

[I]t seems inappropriate to challenge an interlocutor who argues for, say, compatibilism about free will, moral cognitivism or epistemic contextualism, by asking, 'Do you know that?' or 'How do you know?' This provides further evidence for our claim that the contextualist is not asserting the content of her theory. (M&S 2014: 328)

And here I disagree. I hear such challenges often in the settings in question, and on the surface they seem to me no more conversationally inappropriate than are the unhedged statements that they are responses to. So, while M&S conclude, 'This means that . . . our contextualist does not represent herself as knowing the content

of her theory,' I am more inclined toward a treatment on which we do at some level represent ourselves as knowing the controversial things we claim in philosophical settings, but our doing so is excusable, at least to some extent.

In this respect, my approach is somewhat like Williamson's, who views 'lively seminar discussion[s]' as examples where the knowledge norm for assertion is often violated, but the violations are not a big deal, as he expresses in the below passage, which also contains another reason, beyond the consideration discussed in the above paragraph, for why I don't join M&S in denying that the claims in question are assertions:

To rule that the speakers are not making genuine assertions in such situations would be to trivialize the [knowledge] account [of assertion]. In natural languages, the default use of declarative sentences is to make assertions, and the situations at issue are not special enough to cancel the default. Rather, the point is that the knowledge account does not imply that asserting p without knowing p is a terrible crime. We are often quite relaxed about breaches of the rules of the game which we are playing. (Williamson 2000a: 258)

I should note that I don't necessarily join Williamson in charging that excepting the claims in question would 'trivialize' the knowledge account of assertion: Much depends here on what kind of case can be made for taking them to be 'special enough', as Williamson puts it. But I do join him in taking it to be an advantage for a treatment of these claims that it not rule them to be non-assertions. If pointed reactions like 'You don't know that!', or if challenging questions like 'How do you know that?', seemed as out of place in response to the unhedged philosophical claims in question as they are to hedged claims, that would be the kind of consideration that would tempt me to rule with M&S that these claims are not assertions. But I don't find such challenges conversationally inappropriate in that way, so I do classify the claims as flat-out assertions, as Williamson does.

However, though the claims in question don't seem special enough to rule them non-assertions, there does seem to be something *a little* special about them—special enough to make me think a bit more is going on with them than is explained by just Williamson's general 'not all violations of the rule are a big deal' treatment. To return to the contrast presented a few paragraphs above, in settings in which little is at stake, it may not be a high conversational crime to assert to a motorist that 'There's a service station round that corner', when you are nowhere near to knowing that to be the case. And if you are being attacked too vigorously for that conversational misdeed, there would be merit to a defense along the lines of: 'Ease

up on me! What I did was not that horrible.’ But what you did was a conversational misdeed, and, I sense (though I admit this is all very delicate), it is open to criticism in a way (and not just to an extent) that one is not open to criticism for leaving out a hedge when claiming in a heated philosophical debate that, say, free will is incompatible with determinism. It’s this that makes me suspect some kind of *special* license or excuse is at work here. And a key difference between these cases seems to be that the claim you make to the motorist is something that, from her point of view, you really might have known—which makes it badly misleading for you to represent yourself as knowing it by flat-out asserting it. But the philosophical claim is not the kind of thing anyone knows,¹⁴ so there can be some special excuse or license for our engaging, as it seems to me we do, in some kind of *pretense* to know such controversial things in philosophical settings, and thereby flat-out asserting things that we are nowhere near to knowing to be the case.

This pretense of knowledge seems limited in what it licenses. To me, anyway, it does not seem right in such settings to say that you know things that you merely accept. (This *may* be in part because the claim to know some such *p*, because it is longer than is just the flat-out assertion of *p* itself, can’t be justified by how concise and handy it is.) However, as I’ve noted, it does seem, at least to me, to license flat-out asserting these controversial things that we accept, and it also seems to show up in the significance of saying ‘I don’t know’, which in such settings seems to convey that one does not even (choose to) accept the item in question. So, it seems to me that we slip into a mode in which we in certain ways talk as if we know the

¹⁴ Here I come close to the idea that Goldberg explores (though he in the end doesn’t think it works out) and nicely glosses with the nice phrase that in (what I call) ‘knowledge deserts’ we are in conditions of ‘epistemically diminished hope’ (Goldberg 2015: 252; Goldberg takes the idea to be suggested in (Benton 2012: 104-5)). Goldberg explains: ‘The idea would be that, since we all know (and know that we all know) that knowledge is not to be had in these circumstances, and yet despite this we still find that there is a point to trading claims, the result is that “it is understandable why we often enough don’t enforce” the requirement that assertions answer to a robustly epistemic standard in such cases’ (Goldberg 2015: 252). I find this tricky territory, because, though it perhaps *should* be clear to all that we don’t know what we claim in areas of great controversy, we can still to varying extents develop delusions of knowledge here, making it not straightforwardly the case that ‘we all know (and know that we all know) that knowledge is not to be had in these circumstances’. However, in suggesting that these matters aren’t the kind of things anybody is in a position to know (and so perhaps our epistemic hopes *should* be diminished, even where they sadly aren’t), and that this might ground a special license to engage in a pretense (perhaps one we are only semi-aware of engaging in) to know what we do not in fact know, I take myself to be working in the vicinity of the idea Goldberg considers.

controversial positions we accept, even when we don't know them to be true. This allows us to dispense with the various hedges that would otherwise litter too many of our claims in a 'knowledge desert', and to concisely indicate whether we accept a claim by flat-out asserting it if we do, and saying 'I don't know' if we don't.

If I'm right that the latitude we enjoy in this mode of speech does not go so far as to license saying that we know things that we merely accept, one is put in an awkward position in that mode when, concerning some claim that one wants to continue to accept, one is pointedly asked the questions that, unlike M&S, I find appropriate: 'Do you know that?' or 'How do you know that?'. For it is wrong to answer that one does know (if I'm right about our license being limited in that way), but admitting that one does not know, while engaging in the pretense, wrongly conveys that one does not even accept the item in question. I have found the best option in such an awkward situation is to (perhaps temporarily) disengage from the pretense of knowing these matters and 'get real', by answering along the lines of: 'Well, no, I don't know it, of course, but that is my position, and my reasons are ...' This indicates that one is exiting the mode of speech associated with the pretense of knowledge because 'I don't know', which in the mode indicates that the item isn't even part of the position one chooses to accept, is immediately followed by 'but that is my position'.

'Getting real' from time to time in the above way is probably a good idea in any case, even if one isn't put in an awkward dialectical position which pressures one to do so. For if this really is some kind of a pretense of knowledge that we engage in, as I am suggesting, it's a pretense that often gets mistaken for reality, leading philosophers, and others who get carried away by how they feel and talk when operating in knowledge deserts, into (sometimes pathetic) delusions of knowledge. And if this is some kind of pretense, it is also, and relatedly, a pretense that can *feel* very real to us: We are 'method pretenders' (based on the notion of 'method actors'), as it were.

As I speculated several paragraphs above, it may be helpful to philosophical progress that we in a way sincerely feel confident in our acceptance of controversial views. If so, it may prove tricky to avoid overconfidence in philosophy (and in other knowledge deserts) without jeopardizing some of the progress (and perhaps other goods) that such felt confidence might aid. Nevertheless, I am hopeful that recognizing our pretense to knowledge here to be just that, and to be more ready to slip out of our pretense to knowledge and the feelings that come with it, may to

some extent mitigate the sad situation of our often being overconfident blowhards, displaying absurd levels of confidence in our answers to questions we are in no position to be so confident about—without seriously hurting philosophical progress.

Needless to say, I don't have my thinking on these matters very well worked out—and presenting the positions of M&S and of Williamson on either side of mine is not just done for the purpose of clarifying what my position is by contrast to theirs, but also to present readers with some other ways that I remain open to, by which these issues may be handled. This all seemed worth bringing up here, because it seems that working out what to say about assertion (or at least what appears to be assertion) in philosophical settings and other knowledge deserts will also serve to address problems I might face in the special case of articulating my contextualist solution to skepticism. In the meantime, the worry that, as a contextualist responder to skepticism, I face some *special* 'factivity' problem, generated by the invocation of fantastically high epistemic standards, is addressed in section 17 of Chapter 4.