The Ordinary Language Case for Contextualism and the Relevance of Radical Doubt[1]

Michael P. Wolf (Washington and Jefferson College)
Jeremy Randel Koons (Georgetown University in Qatar)

From: *Contemporary Pragmatism* 15 (2018): 66-94.
https://doi.org/10.1163/18758185-01501005

Abstract: Many contextualist accounts in epistemology appeal to ordinary language and everyday practice as grounds for positing a low-standards form of knowledge (knowledge$_L$) that contrasts with high-standards accounts prevalent in epistemology (knowledge$_H$). We compare these arguments to arguments from the height of "ordinary language" philosophy in the mid 20th century and find that all such arguments face great difficulties. However, in Malcolm's metatheoretical justification for his ordinary language methodology, we find a powerful argument for the legitimacy and necessity of knowledge$_L$ (but not of knowledge$_H$). We contend that these appeals to practice do leave us with reasons to accept knowledge$_L$ in the face of the more radical doubts raised by skeptics. We conclude by arguing that by relegating knowledge$_H$ to isolated contexts, the contextualist fails to deal with the skeptical challenge head-on. Knowledge$_H$ and knowledge$_L$ represent competing, incompatible intuitions about knowledge, and we must choose between them. "Radical" doubts that drive most skeptical accounts serve to force confrontations between these intuitions and shape our conception of knowledge accordingly. A fallibilist conception, formed with proper attention to radical doubts, can address the skeptical challenge without illicit appeal to everyday usage.

* * *

What can epistemologists glean from 'common sense' views of knowledge, particularly with respect to epistemological skepticism? Many efforts to appeal to some forms of common sense on these questions historically ran through 'ordinary language philosophy', while a different set of accounts have more recently been offered by epistemic contextualists. There is a surprising confluence between these two schools of thought. In both cases, fundamental assumptions about how to offer an account of knowledge proceed by appeals to everyday practice, and in both cases, there is the promise of answering the skeptic while preserving most knowledge ascriptions. But common sense is more fraught than it first seems, and not all appeals to everyday practice succeed in shoring up our theoretical accounts. We begin by discussing contextualist approaches to semantics and epistemology, then critique their appeals to common sense and ordinary language. We close with our own analysis of the reasons for a fallibilist account of knowledge that addresses some of the shortcomings of contextualist and ordinary language accounts.

1. Epistemic Contextualism

We seem to have competing (and prima facie incompatible) intuitions about knowledge, intuitions that support skepticism and others that support the legitimacy of everyday knowledge attributions. One set of intuitions seems to support a high-standards, even infallibilist, conception of knowledge according to which any possibility of

[1*] Our thanks to Gregg Osborne and Jeremy Morris for helpful comments on earlier drafts.

error is incompatible with knowledge. According to this Cartesian conception of knowledge (knowledge$_H$), knowledge requires certainty. The other set of intuitions exploited by contextualism supports a low-standards, fallibilist conception of knowledge (knowledge$_L$), according to which knowledge is compatible with the possibility of error. Much mainstream epistemological debate is between contextualists, who try to reconcile these intuitions, and invariantists, who must settle on a single conception of knowledge.[2] Cohen (1998, 2000) DeRose (1995, 2002) and Lewis (1996) have set the agenda for most epistemic contextualists.

It is difficult to reconcile these two sets of intuitions, and skeptical arguments have stubbornly resisted refutation; we won't rehearse the various attempted refutations in detail. What concerns us is the novel approach the contextualist has taken toward reconciling these intuitions. The contextualist begins with a contextualist account of the semantics of 'knows'. Plausibly, many terms in everyday use have contextually-sensitive semantics, e.g. indexicals and pronouns. According to some contextualist accounts of terms such as 'tall,' 'flat,' 'large,' 'small,' and so on, a sentence employing one of these terms can express different propositions in different contexts. Thus, writes John MacFarlane (2007), when he says of his niece, "Chiara is tall," this sentence expresses the proposition "Chiara is taller than the average 7-year-old" and is true. But when her basketball coach utters the same sentence in the context of discussing who should play which position on the team, the sentence expresses the proposition "Chiara is significantly taller than the average player on the team," and is false.[3]

Similarly, for the epistemic contextualist, the content of a sentence of the form "S knows that P" can change with context; it can express a different proposition, with different truth conditions. Thus, the sentences "I don't know that I am not a brain in a vat" and "I know that I am typing on a computer" only entail contradiction if we assume an invariantist conception of knowledge. But if we assume that the sentences attribute knowledge according to different standards–one a high standard appropriate to skeptical contexts, and the other a low standard appropriate to everyday contexts–the two knowledge attributions don't conflict anymore than the two utterances of "Chiara is tall" conflict with each other. For the sake of clarity, in the future when we are discussing attributions of high-standards and low-standards knowledge, we will speak of attributing knowledge$_H$ and knowledge$_L$. A

[2] A wild card in this debate is subject-sensitive invariantism (SSI), as defended by Hawthorne (2004) and Stanley (2005). Wherever one comes down on this question, our arguments here bear on features of contextualism that don't figure prominently in SSI.

[3] MacFarlane, incidentally, rejects contextualist accounts of such terms in favor of what he calls 'semantic minimalism.'

contextualist solution to the the problem of radical skepticism has been offered by, among others, DeRose (1995), Cohen (1988 and 1999), Neta (2003).

Two comments are necessary at this point to head off misunderstanding. First, we realize that this approach threatens to oversimplify the contextualist theory. According to the epistemic contextualist, we don't make just two kinds of knowledge attributions–attributions of knowledge$_H$ and knowledge$_L$. There are many (perhaps uncountably many) standards, dependent on the various contextual features of the subject and/or attributor of knowledge. Nevertheless, various contexts can be described as 'skeptical contexts'. All such contexts employ some version of a high-standards skeptical hypothesis, which we purportedly cannot refute, which defeats all knowledge-attributions. Thus, whatever minor variations in epistemic standards might be at play in such contexts, we shall say that all such contexts are knowledge$_H$ contexts. Similarly, other contexts will, though perhaps involving a variety of different standards, be fallibilist in nature–they will admit of the possibility of knowledge of an everyday sort that doesn't require incorrigibility or infallibility. In these contexts, we can say that knowledge$_L$ is being attributed or denied. Even within everyday situations, the contextualist will argue that there are situations that demand higher (but non-skeptical) standards, e.g. Cohen's (1999) Airport Case. We might call these sorts of contexts knowledge$_M$ cases. We will discuss these sorts of cases later in the essay.

Second: although we will, for the sake of clarity, often talk about knowledge$_H$ and knowledge$_L$, we are cognizant of the fact that for most contextualists, theirs is a theory not of knowledge per se, but of the semantics of 'knows' and cognate terms. Thus, they are giving not a theory of knowledge, but a theory of when it is appropriate to attribute knowledge. We ask for some forbearance when we use terms like 'knowledge$_H$' and 'knowledge$_L$' in connection with the contextualist theory. Although contextualism is a theory about the semantics of 'know' and cognate terms, it is needlessly cumbersome to avoid object-language talk in writing about contextualism. So the reader is to bear in mind that even when we make claims about, e.g., knowledge$_L$ and knowledge$_H$, these claims can be translated into claims about sentences of the form "S knows$_L$ that P" and "S knows$_H$ that P," and the contexts in which such sentences can be true.

2. Arguments from Ordinary Language

Contextualists appeal to our intuitions about ordinary cases to convince us that the proposition expressed

when one attributes knowledge varies with context.[4] In ordinary contexts, such attributions impose lower standards, and knowledge attributions can more easily be true. In so-called "philosophical" contexts, the standards are higher, and it may seldom (or never) be correct to attribute knowledge in such contexts. The case that we do and should make such shifts begins with an appeal to ordinary usage.

> The best grounds for accepting contextualism concerning knowledge attributions come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk: what ordinary speakers will count as 'knowledge' in some non-philosophical contexts they will deny is such in others. (DeRose, 2005, p. 172)

DeRose qualifies that passage a bit by adding that speakers say such things "with propriety," but the grounds for this are themselves appeals to ordinary language usage (other speakers accept their doing these things). It might be more precise to say that epistemic contextualism is put forward as a theory, and that its best defense is that it best explains our 'common-sense' epistemic intuitions. However, this still takes those intuitions and their expressions in ordinary language as settled points to which any account must comport itself. Critics might challenge contextualists in a number of ways here.[5] However, such critics typically take ordinary language to establish something, even if it does not resolve epistemological problems like skepticism. We doubt that the way in which ordinary language is called upon here does all of our commitments justice.[6]

Our use of a word reveals more than our semantic commitments. Often (perhaps always), the linguistic norms governing the use of a word embody certain theoretical commitments, and theoretical commitments are always contestable. The fact that ordinary language enshrines a particular set of theoretical commitments in no way legitimizes those commitments, or shields them from critical scrutiny. Consider a contemporary example. Many opponents of gay marriage advance an essentially semantic argument against the notion of same-sex marriage. They argue that 'marriage' just *means* 'a union between a man and a woman', and so it is semantically impossible for

[4] Most literature on contextualism distinguishes between semantic contextualism and epistemic or substantive contextualism. Most contextualists defend the former sort of view, and (as should be apparent from our discussion so far) much of what we say is addressed toward this dominant strain of contextualism. However, our defense of a low-standards conception of knowledge remains viable whether one is talking about what proposition is expressed by a knowledge-attributor (as the semantic contextualist does) or whether one is "concerned with making substantive claims about knowledge or justification itself" (Rysiew 2011, p. 6), as is the epistemic or substantive contextualist.

[5] E.g., Kornblith (2000), Sosa (2000), Stanley (2004), and Brueckner (2005).

[6] We will not dwell on the fact that the goal of the original practitioners of OLP was not to legitimize skeptical intuitions, but to show how fundamentally misguided they were.

there to be a marriage between two individuals of the same sex.[7] This example differs in detail from the contextualist example, as in this example, what is at issue is the meaning of a contested term, rather than what proposition is expressed when the term is used in a particular context. But this difference isn't relevant to the question at hand: can appeal to ordinary usage settle moral, epistemological, and other philosophical disputes? This argument is pre-Quinean; it assumes that there is a fundamental cleavage between meaning and theory, and that theoretical considerations can never compel a meaning change. But such a sharp divide between meaning and theory is surely not tenable. If a resident of the US in the early 1800s were to argue that blacks cannot be citizens because 'citizen' just *means* (in part) 'white person', then we should say that there are powerful *moral* reasons for adjusting the definition of 'citizen', or jettisoning the old concept in favor of a new one. Thus, there is overlap between moral and semantic considerations; the former can bear on the latter, and the latter aren't wholly separate and inviolable.

Similar arguments have appeared in more traditional philosophical territory. For example, Stace (1952) holds that proper ascription of free will is settled by how the term is used in ordinary language. But again, ascribing free will to someone has normative consequences for how they are treated, and whether this theoretical structure of freedom and moral responsibility can survive critical scrutiny isn't ipso facto resolved by a consensus among speakers. As these examples demonstrate, the method of ordinary language can show us how a word is used; it can show us the ordinary meaning of a word. It cannot, however, determine whether this ordinary meaning can bear theoretical scrutiny. For the meaning of a word is tied up with various theoretical commitments, and theoretical commitments are inherently contestable. Thus, the method of ordinary language philosophy may tell us that when we say "S knows that P" in an ordinary context, we are expressing the proposition, "S knows_L that P;" this may in turn imply that S is entitled to ignore certain remote possibilities. But appeals to ordinary language don't explain where the *entitlement* to ignore these remote possibilities comes from, except by noting that people commonly do so. *Appeal to ordinary usage alone is not sufficient to justify a theoretical commitment*. This is the lesson of our examples regarding the terms 'citizen', 'marriage', and 'free will'; and it is a lesson that applies equally to the word 'knows' and its cognates.

One might argue that we aren't being fair to contextualists, and the specific way that they practice ordinary

[7] The argument isn't altered if the locus of dispute is shifted from meaning to reference, i.e., if one argues that the word 'marriage' designates these same-sex marriages, and that this fact of designation is simply a semantic fact, established by ordinary usage. One can still argue, based on moral or other considerations, that this concept of marriage ought to be discarded in favor of one that designates more broadly, or that ordinary usage should be amended so that other sorts of arrangements are also designated, etc.

language philosophy. Some practitioners of OLP are happy to accept a simple inference from "S is the ordinary usage of concept C" to "S is the *correct* usage of concept C," but surely contextualists *are* more sophisticated than this. Nevertheless, contextualists do consider our intuitions about the proper usage of a term to be the strongest evidence we have for the truth of certain kinds of knowledge attributions. Consider DeRose, who has considered these issues more carefully than most. DeRose's OLP is based on a commitment to what he (borrowing from Brian Weatherson) calls the 'methodology of the straightforward', which "takes very seriously the simple positive and negative claims speakers make utilizing the piece of language being studied, and puts a very high priority on making those natural and appropriate straightforward uses come out true…" (DeRose 2005, p. 192).

For DeRose, while our intuitions and ordinary (proper) usage don't directly entail the correctness of a particular employment of a theoretical term, they are are strongest evidence for it, and represent powerful "mutually reinforcing strands of evidence." But DeRose argues further, that if people use a term like 'knows' with propriety, then this, in turn, is strong evidence that such claims are *true*:

> [S]peakers do in fact use 'knows' in the way described, and appropriately so – they will in fact, and with apparent propriety, ascribe 'knowledge' in situations like LOW, yet will deny 'knowledge' when they find themselves in conversational circumstances like HIGH. This supports the premises that both of the imagined claims are true, since generally (though there are some exceptions), one cannot properly claim something that from one's own point of view (given one's beliefs about the underlying matters of fact relevant to the claims in question) is false. So, since the contextualist's cases do not involve speakers who are involved in some mistaken belief about a relevant underlying matter of fact, there is good reason to think that their claims, which are made with perfect propriety, are true, and it's a bad strike against a semantic theory if it rules these claims to be false… (2005, p. 173).

Thus, if speakers use a term (like 'knows$_L$') with propriety, and based on no false beliefs, then we have good reason to think that what is said is true. So DeRose has offered us, in effect, an argument starting from the premise that speakers of English use 'knows$_L$' with propriety, and draws as a conclusion that attributions of knowledge$_L$ are (presumably) generally true. Other contextualists, despite their sophistication, make essentially similar arguments: if 'knows' is ordinarily used in a particular way, and our intuitions support this as a proper use of the language, then this provides decisive evidence that such uses are correct, and such knowledge attributions are, by and large, true.[8]

We deny that the enshrinement of these attribution-types in ordinary usage somehow legitimizes them. This argumentative gambit corresponds roughly with the final stage in DeRose's argument where he claims that the knowledge attributors in his cases "with propriety use the claims in question" and proceeds via "the presumption

[8] See Travis (1989), Klein (2000), Kompa (2002), Neta (2002, 2003), Ludlow (2005), Montminy (2007).

that what is properly said is true" (2005, pp. 173, 181) to the conclusion that the knowledge claims enshrined in ordinary usage are (probably), by and large, *true*. Again, merely showing that ordinary usage enshrines multiple, contextually-sensitive uses of the knowledge concept (or that we are changing the subject in philosophical contexts) does not in itself *vindicate* any of these uses.

Before beginning our positive argument for the legitimacy of knowledge$_L$ attributions, there are two points that need clearing up. First, it is important to distinguish our criticism of contextualism from earlier, prominent criticisms of contextualism. A common criticism of contextualism has been that because it is a semantic or metalinguistic thesis, it does not address the skeptic's concerns at all. For the skeptic won't merely concede that while knowledge-attributions aren't true in high-standards contexts, they are (or can be) true in low-standards contexts. On the contrary, the skeptic is challenging the possibility of any sort of knowledge; the skeptic is arguing that even our *ordinary* standards fail to be satisfied (say, because we have *no evidence at all* that we aren't brains in vats). Thus, Kent Bach writes, "the contextualist's attempt to marginalize skeptical arguments by restricting them to skepistemic contexts [contexts where skeptical standards prevail] ignores the fact that skepticism denies that we have knowledge even by *ordinary* standards" (2005, p. 68). Virtually identical points are made by Sosa (2000, p. 6), Kornblith (2000, p. 27), and Feldman (2001, esp. section 3).

The argument we are making is importantly different (although we will revisit an argument similar to the above one when we argue for the 'explosiveness' of radical doubt later in the paper). We argue that having a low-stakes conception of knowledge in the first place is in need of defense (and the contextualist's ordinary language defense is, as we saw, unsound). Even if we can meet the lower standards that are presupposed by attributions of knowledge$_L$, why should we think that a doxastic state with only a moderately strong epistemic standing should count as *knowledge*–should count as a state we are epistemically entitled to hold? The above critics of contextualism point out that according to the skeptic, these lower standards are never met; we ask the question, "Even if these standards *are* met, should we count this achievement as *knowledge*?" This is essentially a question of which standard we ought to adopt. We will argue later that far from providing a happy reconciliation between knowledge$_H$ and knowledge$_L$, we are forced into a confrontation between the two, and the epistemologist must decide which conception of knowledge to embrace. For, as we will argue, even if knowledge does display some contextualist features, there is no room for both a low-stakes, fallibilist conception of knowledge and a high-stakes, infallibilist conception. Knowledge$_L$ and knowledge$_H$ cannot happily coexist, and the preponderance of reasons fall

on the side of keeping knowledge$_L$ and jettisoning knowledge$_H$.

Second: one might argue that epistemic contextualists are merely natural language semanticists, and as such are merely engaged in a descriptive endeavor. They don't need to 'vindicate' anyone's use of the word 'knows', since epistemic contextualism is merely a descriptive linguistic view with interesting philosophical consequences.[9] This objection throws the epistemic contextualist on the horns of a dilemma. On the one hand, she can say that she is trying to solve the skeptical paradox, which manifestly motivates many (perhaps most) contextualists. But as we have repeatedly been urging, the contextualist cannot justify ignoring skeptical doubts merely by pointing out that we do, in fact, do so in ordinary contexts. The fact that the ordinary conception of knowledge we use expresses a certain low-standards content in ordinary contexts doesn't show that this practice can be defended.

On the other hand, the contextualist can say that theirs is merely a descriptive semantic theory, and not one that is intended to solve skeptical problems. This is unsatisfactory. To merely show that 'knows' expresses different propositions (ranging from low-standards in ordinary contexts to high-standards in skeptical contexts), but not to claim the legitimacy of any of these ways of speaking, is unsatisfactory for two reasons. First, it does not add anything of philosophical importance to the discussion. We already knew that we were inclined both to ascribe knowledge and withhold knowledge ascriptions. If contextuajonelism can't tell us which way of speaking is correct–or whether we are ever correct in ascribing knowledge to ourselves or others–then it's not a philosophically-interesting thesis. Second, to make a knowledge-attribution is to make a normative assessment—it is to attribute entitlement to a belief or proposition. For the contextualist to say, "I am merely saying that there are different kinds of knowledge ascriptions that we make, which express a variety of different propositions, embodying different standards," is to evade the central question—are any of these attributions true? Does anyone ever achieve the entitlement that is attributed by these claims? Thus, what we would need here is an argument for the *legitimacy* of low stakes knowledge, even in the face of concerns that often lead people to high stakes positions. Even better would be one that legitimized low stakes knowledge in general. We do wish to argue for the legitimacy of knowledge$_L$ and cognate notions. As it turns out, an argument for this legitimacy isn't far away.

3. Malcolm and Ordinary Usage

We will return to our discussion of contextualism shortly. But first, we want to argue that although the

[9] An anonymous reviewer raised this objection.

ordinary language argument deployed by the contextualists fails, there *is* a viable fallibilist conception of knowledge (i.e., knowledge$_L$) and justification. Consider the following statements:

> Kansas is flat.
> This pool table isn't flat.

These sentences might seem paradoxical, as any pool table is flatter than Kansas (and the contextualist draws aid and comfort from the fact that a structurally similar paradox can be created using the word 'knows'[10]), but we immediately recognize that the word 'flat' is being used in two different contexts, and thus has two different contextually-determined standards. Thus, the two sentences above are rendered:

> Kansas is flat$_L$ (i.e., it has no mountains, the terrain isn't hilly)
> This pool table isn't flat$_H$ (i.e., the felt is wrinkled)

reflecting a low and a high standard of flatness.

Why not respect ordinary usage in this? Clearly, there is great utility in having the term 'flat' express different contents in different contexts. We would not be made better-off by adopting a single universal standard of flatness, by using an invariantist semantics for 'flat'. Clearly, there are distinct, contextually-determined standards for the word 'flat', and these distinct standards are legitimate. But is this conclusion yielded by the ordinary language methodology, or by considerations distinct from it?

There is a good reason for keeping and using 'flat' to express the content we shall call 'flat$_L$', but that reason does not stem directly from its enshrinement in ordinary usage. Rather, the reason stems from the utility of keeping this content. In particular contexts, this content serves our linguistic and communicative interests, and it would be difficult to achieve them without this content. It is helpful, when describing terrain, to have a variety of low-standard versions of 'flat' available to us. For pool tables, or high-tech components, or theories of geometry, other standards may hold sway. They aid us in achieving our interests, including our communicative goals. Thus, various terms and contents are enshrined in ordinary usage in part because of their utility.[11]

---

[10] The above sentences can be made paradoxical by the insertion of a plausible auxiliary premise such as "If A is flat, and B is flatter than A, then B is flat." This premise creates a paradox precisely because it uses a context-invariant sense of 'flat' to connect two premises, each of which uses a context-sensitive notion of 'flat' (and each in a different context). A structurally similar move is how the contextualist claims skeptical paradoxes are created using the word 'knows'.

[11] The fact that such terms are embedded in ordinary usage contributes to their utility, but in considering the defensibility of a term (such as 'free', 'knows', 'flat$_L$', etc.) the common usage of the term isn't relevant independent of how it contributes to the utility of the term. If one considers ordinary usage as a reason for maintaining usage of a term, it is only insofar as common usage of a term contributes to the term's utility by contributing to the term's ability to advance our linguistic goals (i.e., by making the term widely-understood in the community).

This latter fact was recognized by Malcolm, and served as a metaphilosophical justification for his own ordinary language methodology. But crucially, it isn't merely ordinary usage that justifies continued usage. That is fallacious reasoning, on par with *argumentum ad populum*. We are justified in continuing to use, to take Malcolm's example, 'clear' because there are things we cannot express if we jettison this term and retain only the term 'vague.' And so instead of merely appealing to ordinary usage, the contextualist needs to argue that important epistemic needs are fulfilled by retaining both 'knows_L' and 'knows_H', and that we cannot make do with a single, invariantist conception of 'knows'.

As it turns out, though, no such argument can succeed. On the contrary, there are important communicative needs that can only be met by a fallibilist conception of knowledge, and not by an infallibilist conception of knowledge. These communicative needs can only be met by the retention of 'knows_L', and cannot be met if we retain only 'knows_H'. (And, we will see, we cannot keep both.)

If the only epistemic tools we have are knowledge_H and related notions, then communication becomes impossible. We run up against an important practical facet of knowledge ascription, described by Hawthorne[12]:

> Plausibly, when someone asserts that *p*, she conveys that she knows that *p*…The practice of assertion is constituted by the rule/requirement that one assert something only if she knows it. Thus, if someone asserts *p*, it is proper to criticize that person if she does not know that *p*. (2004, p. 23)

Note, for example, the oddity of the following conversation:

> Liz: "Is that the 54C bus that is coming?"
> Flynn: "Yes, it is. And I don't know whether it is or not."

The oddness of the conversion is explained thus: when one makes an assertion, one is expressing a commitment to the proposition, as well as entitlement to the proposition. Sellars (1956/1997) puts the point in terms of endorsement; key to his discussion of 'looks-talk' is that what distinguishes looks-talk from ordinary assertion is primarily the absence of this dimension of endorsement: to say that something *looks* red is to withhold endorsement of the claim that *it is red*. His point there is largely about epistemic priority, but it may inform discussion here. One may commit to a proposition without demanding that others endorse it, and perhaps the skeptic imagines our doing so on an everyday basis, thereby obviating the need to talk of knowledge here.

[12] The idea that knowledge is the norm of assertion has been defended, in various versions, by Williamson (2000) and DeRose (2002), and can be traced back at least as far as Unger (1975).

However, this misses the more urgent point we are making. If we accept high standards for knowledge and expect anyone who makes an assertion to meet these, then we will virtually never have the standing to make any assertion. It does the skeptic no good to say that we might retreat to accepting that P, or believing that P. Each of these recapitulates the same problem: part of the constitution of each of these standings is an endorsement that high standards would preclude us from making. Thus, adopting knowledge$_H$ as our standard would not imply a minor diminution of the set of assertions we may make (subtracting only most knowledge ascriptions), but a more fundamental degradation of communicative discourse.

Suppose Smith and Jones are running an experiment, and Smith asks Jones whether the bacterial colony in question has developed resistance to the toxin introduced into its environment. Now imagine that Jones has become a high-standards invariantist about knowledge, and thus acknowledges only the existence of knowledge$_H$. How shall Jones reply? He cannot assert "Yes, it has," for he does not know that it has—at least by his own lights, for he only acknowledges the legitimacy of knowledge$_H$. Nor can he assert, "No, it has not," and for the same reason. It seems that the only response Jones can give to Smith's question is "I don't know$_H$." Jones might retreat to saying that he believes that the colony has developed resistance, and report this only as a matter of his own conviction and commitment, but this just isn't what Smith was asking. Smith was not interested in Jones's commitments or psychological states themselves, he was asking for a report on the colony. If knowledge$_H$ forces us to make such retreats globally (or nearly so), then it forces us to withhold endorsements and licenses for action across the board, gutting any promise of communicative action.

A critic of argument might reply that the epistemic norm of assertion needn't be as strong as knowledge, and that we can still make assertions and engage in communication even if knowledge$_H$ is the only type of knowledge there is. We often make statements of belief with greater or lesser degrees of confidence. Some of them are delivered with very little confidence at all and tremendous hedging of our bets ("I believe the world will find more effective replacements for fossil fuels some day") and others may be delivered with much higher confidence, only a formal commitment to the possibility of error, and no real further effort devoted to checking them ("I believe that 37 is prime"). Thus, our bus-stop example from above might read:

> Liz: "Is that the 54C bus that's coming?"
> Flynn: "I believe it is, and I believe so as strongly as I reasonably can. The 54C is the only bus that comes to this stop, it comes at almost exactly this time every day, I have good vision and can see the sign on the bus from here, and nothing I can discern is interfering with my perception. There is some chance that I cannot eliminate that I am subject to some illusion, hallucination, or that I am merely a brain in a vat. But those are possibilities so far divorced from the ways I can engage with

> things that I cannot see how they can have any practical bearing for us. Let's proceed on the assumption that those possibilities do not obtain unless something compels us to change our approach."

Hasn't Flynn just made a legitimate assertion, which yet doesn't abide by the epistemic norm of assertion (ENA)? There are two things to say here. First: we agree that Flynn could say something like this. We just think that by a reasonable interpretation of our epistemic practices, he thereby knows that the bus is the 54C and should just assert it. Those 'low stakes' standards conflict with 'high stakes' standards, but what is at issue here is precisely which set of standards (and which corresponding conception of knowledge) is of any use in our language. Knowledge_{L} allows Flynn to actually make an assertion and communicate information.

But (and here is the second point) high-standards knowledge does not allow Flynn to make any assertion or communicate any information. The skeptical argument works just as well to undermine the possibility of justification as well as knowledge. If the possibility that Flynn is a brain in a vat, or deceived by an evil demon is genuinely in play, then not only doesn't he know_{H} that the bus is the 54C, but how can he even be justified_{H} in believing that it is the 54C? And if Flynn lacks entitlement to his claim, he cannot make an assertion–he lacks grounds for endorsing the proposition that it is the 54C (or that it is probably the 54C, or likely the 54C, or the 54C to the best of his evidence, since he has no evidence [i.e. $\text{evidence}_{H,}$],...). Thus, even if the norm of assertion is weakened from knowledge, knows_{H} still makes it impossible to assert anything at all, because one cannot ever be in a position (have the proper entitlement) to make an assertion. So retreating from knowledge as the norm of assertion doesn't solve the problem–one would have to say, instead, that assertion is governed by no norm of entitlement *at all*, which is wildly implausible.

We said above that there were two replies that the critic might make at this point. We have just discussed the first (the argument that communication is still possible if we only have knowledge_{H}). The second reply is simply to deny the ENA. For example, Jennifer Lackey (2007) has presented a number of widely-discussed counterexamples to the ENA, particularly the example of 'selfless assertion.' In a typical example, Lackey describes Stella, a devout Christian creationist who acknowledges that the evidence for evolution is overwhelming and that her own belief in creationism is based on faith rather than reason. When Stella is teaching her fourth-grade science class, she asserts "Modern day Homo sapiens evolved from Homo erectus," which she doesn't believe (and ergo doesn't know), because she feels she must present to the students that which is best-supported by the evidence.

It would be easy to deny that Stella is really asserting the sentence in question (as she lacks the relevant

commitment). But we needn't even engage this argument. Lackey's examples of selfless assertion involve a person who lacks the relevant belief, *but who possesses the corresponding entitlement*. Our entire argument has been that limiting our epistemic resources to knowledge$_H$ (or justification$_H$) would make assertion impossible, because it would make impossible this element of entitlement that is crucial to assertion.

Other supposed counterexamples to ENA seem no more damaging to our case. Cappelen (2011) provides a number of counterexamples to the ENA. Some of them (such as 'political bullshit', statements by political figures which "everyone knows to be false") stretch the definition of assertion. The implausibility here is due to the link between assertion and testimony: assertion (even selfless assertion) is a means of conveying (or attempting to convey) information and entitlement, which is why assertion typically presupposes entitlement to its contents in the first place. But a statement to which one is neither committed nor entitled hardly qualifies as an assertion.

But even if Cappelen's counterexamples are genuine, and there are instances of genuine assertion without entitlement (or without commitment, which seems to be the more common type of example), what does this show? Does the fact that we can make *some* assertions without entitlement mean that we can do away with the notion of entitlement altogether? This is a preposterous non-sequitur. There may well be exceptional kinds of assertion (Cappelen's example of football fans cheering "We will win the championship!") which one can make without proper entitlement.[13] But when we think of one of assertion's central and most critical functions–conveying information among people, and coordinating action–we see that this particular function cannot take place without a viable notion of entitlement. One simply cannot be indifferent to whether one is entitled to the claims one makes, when one endeavors to convey information to one's fellow rational agents. To do so is to violate fundamental principles of communication. When one undertakes to convey information to another person, one undertakes to convey to another how the world is. And as a rational agent, one must care about the way the world really is. But as we saw, if the only notions of entitlement we have are knowledge$_H$ and justification$_H$, one lacks the standing to assert any information. Any cautious assertions actually made under such conditions represent a surreptitious utilization of knowledge$_L$ and its cognates, without acknowledging the use of such.

There is one final objection we will consider: why are our only options here a knows$_L$ and the skeptic's knows$_H$? Can't discourse be governed by another sort of high-standards knowledge, the sort that is operative in

[13] For the record, we don't buy Cappelen's argument that these are genuine assertions at all. The supposition leading this paragraph is *entirely* hypothetical.

DeRose's high standards Bank Case or Cohen's Airport Case? Assertion appears possible if we adopt this standard of knowledge–call it 'knowledge$_M$'. We don't object to the claim that knowledge might be context sensitive in some ways–that in some contexts, knowledge$_L$ might be in play, and in other contexts (say) knowledge$_M$ might be in play. Two points need to be made here. First, both of these non-skeptical conceptions of knowledge need to be vindicated in the first place, and mere enshrinement in ordinary practice does not suffice. Let us grant that assertion is possible on both; what we now need to do is establish that one or the other is a viable conception of knowledge. (For the sake of simplicity, we are going to argue that knowledge$_L$ is viable, and leave it open that knowledge$_M$ is also legitimate operative in some circumstances.) Second, as we will argue later (in section 4), we ultimately have to make a choice–a choice between an infallibilist conception of knowledge (knowledge$_H$), and a lower-standards, fallibilist conception of knowledge (whether this be an invariantist conception of knowledge like knowledge$_L$, or a contextualist but fallibilist standard of knowledge varying between knowledge$_L$ and knowledge$_M$). Either way, the dialectical strategy of this paper is dictated by this choice: we must choose between fallibilism and infallibilism, and so we must argue that there are good reasons not to choose an infallibilist standard of knowledge, and there are good and sufficient reasons to choose a fallibilist standard of knowledge. Again, for the purposes of this paper, we are arguing for the legitimacy of knowledge$_L$, but don't rule out the possibility that the content of knowledge-attributions may vary (within fallibilist parameters) depending on context.

So we have seen that if the only epistemic notion we have is knowledge$_H$, communication and assertion are impossible. We have also seen that the contextualist's ordinary language argument provides us with no good reason for retaining knowledge$_L$. What is to be done?

Malcolm, the ordinary language philosopher, can help us. Consider Malcolm's example: the thesis that all words are vague. What does this mean, and why would a philosopher make this claim? Malcolm notes that with any word, there are "undecidable cases" where "the question is raised as to whether the word applies or not" (1968, p. 121). The only difference between clear words and vague words is in the relevant number of undecidable cases. But "the difference between a large number of undecidable cases and a small number is only a difference of *degree*! He is, therefore, tempted to say that all words are really vague" (1968, p. 121).

"But," Malcolm asks, "why should not the use of the words 'vague' and 'clear,' in ordinary language, simply serve to call attention to these differences of degree?" (1968, p. 121) As Malcolm points out, there are words

in ordinary language that function as opposites ('large' and 'small', 'animate' and 'inanimate', 'vague' and 'clear', etc.). One member of the pair only has meaning when contrasted with the other member, so if we drop the word 'clear' from our language and denote everything as vague, "we should have gained nothing whatever."

> [V]agueness was *contrasted* with clearness. In the revised language vagueness could be contrasted with nothing. The word 'vague' would simply be dropped as a useless word. And we should be compelled to adopt into the revised language a new pair of words with which to express the same distinctions formerly expressed by the words 'clear' and 'vague.' (1968, p. 122)

A similar point can be made about the word 'know_L' as it appears in ordinary language. Clearly, 'know_H' and 'know_L' aren't contraries in the same way that 'vague' and 'clear' are; and the contextualist isn't claiming that the meaning of 'know' changes with context, only that the content expressed by use of 'know' changes. But an argument similar to Malcolm's can be made with respect to these different contents expressed by 'know'. Suppose that we (like Malcolm's philosopher) agree to jettison knowledge_L, and make do with knowledge_H. This sort of infallibilism makes knowledge all but unobtainable. But if we retain only the content knowledge_H, then (as with 'vague' and 'clear'), an important distinction gets lost. We cannot distinguish between the person who doesn't have knowledge because he is guessing, or has consulted an astrologer, or is too lazy to investigate his claims, and the person who has diligently gathered evidence but doesn't have knowledge because he cannot rule out some outlandish skeptical hypothesis (such as the possibility of being a brain in a vat).

Consider Smith and Jones, and their standing toward proposition P:

P: The sun is made of hydrogen.

Suppose Smith has not investigated the truth of P, and does not have P on the word of any authority, but merely believes P on the basis of tossing a dart at a periodic table of elements. Jones, on the other hand, has read many scientific papers on the subject, and has personally used a spectroscope to examine the light emitted by the sun, and by examining the solar spectrum has been able to determine the gaseous composition of the sun. If the only conception of knowledge we have is knowledge_H, we must say that neither Smith nor Jones knows_H that-P. But a crucial difference is being elided here. As with Malcolm's example of 'vague' and 'clear', our philosophical "progress" has actually impaired our ability to make an important distinction: the distinction between the epistemic accomplishments of Smith and Jones. It is true that neither has knowledge_H. But there is a crucial difference between Smith's belief in P and Jones's belief in P, and nothing is gained by assimilating one to the other. On the contrary, much is lost if we lose our ability to state what separates Smith and Jones. Of course, a distinction may serve some

theoretical purpose even if it is never realized. 'Flat' in a strict Euclidian sense, or 'frictionless' in a physicist's model are features that belong to nothing in this possible world, perhaps in none at all, but the distinctions serve useful theoretical roles as idealizations. But here, our concerns about assertion and endorsement must come to the fore. We need a distinction that separates Smith from Jones and licenses actual assertions in real, flesh and blood cases. Some forms of knowledge$_L$ will do that, but not knowledge$_H$.

It may seem as though we are defending the methodology of ordinary language philosophy, and by extension its application by contextualists to the semantics of 'knows.' However, this isn't the case. What makes this part of Malcolm's argument particularly effective is that he is appealing to the *utility* of a particular distinction rather than appealing to its enshrinement in ordinary language. To argue for the legitimacy of knows$_L$, we must argue for the utility and theoretical defensibility of this content. We have already taken several steps in this direction; we shall now take several more.

So there is a pragmatic justification for adopting and maintaining knows$_L$, but does this practice survive theoretical scrutiny? We cannot defend fallibilism in detail, but let us address a natural objection that arises here: "It might be *useful* to employing a fallibilist conception of knowledge, but what *justifies* you in ignoring these remote possibilities? After all, you are the one who insists that the epistemic concepts we employ should be able to withstand theoretical scrutiny." Two things need to be said in reply to this objection. First, knows$_L$ ought to be retained because it says something we need to be able to say, something we cannot say using only 'knows$_H$': namely, that (say) Smith's epistemic achievements have significantly surpassed those of Jones, and Smith has achieved a level of epistemic accomplishment that (for our present discursive purposes) deserves recognition by awarding him entitlement to certain claims. Thus, as we have argued above, knows$_L$ is a content that allows us to express a genuine epistemic distinction and thus has theoretical legitimacy.

Second, this objection merely represents the intrusion of intuitions regarding knows$_H$. We understand the source of these intuitions regarding knows$_H$. As we noted earlier, contextualism gains plausibility because it recognizes that we have both infallibilist intuitions (supporting the deployment of a high-standards conception of knowledge) and fallibilist intuitions (supporting the deployment of knows$_L$). Ultimately, the infallibilist intuitions (those supporting the deployment of knows$_H$) must be abandoned in favor of fallibilist ones (those supporting the deployment of the more pragmatically useful conception of knows$_L$). If we can show that infallibilism should be dispensed with in favor of fallibilism, then we can show that knows$_H$ should be discarded in favor of knows$_L$, and

that the above objection—representing as it does an objection from the perspective of $knows_H$/infallibilism against $knows_L$/fallibilism—can be set aside. An important lesson we have learned is that while our Malcolmian argument teaches us the importance of $knows_L$ as an epistemic standard, and as a theoretical construct which all of our communication and action presupposes, no such Malcolmian argument has been made in favor of $knows_H$. On the contrary, we saw that $knows_H$ undercuts the very possibility of knowledge, communication, action, and human community. What lesson should we take away from this?

Recall the metaphilosophy behind Malcolm's ordinary language method. The mere fact of a linguistic practice's enshrinement in ordinary language could not justify the continued inclusion of that practice in our language. There are further questions to ask: does this practice advance the communicative goals of the linguistic community? Do the theoretical commitments implicit in this practice withstand critical scrutiny? As we argued earlier, adopting $knows_H$ simply doesn't pass the test of communicative utility. It adds nothing to the language; it doesn't allow us to express any additional useful content (and indeed, subtracts substantially from our ability to express content). Thus, *properly understood, the ordinary language method does not support both $knows_L$ and $knows_H$–it only supports the retention of $knows_L$ within the language*.

4. Interests and Shifting Standards

We have argued above that the contextualist's appeal to ordinary usage in justifying shifting standards isn't justification-conferring. We have seen, however, that a particular linguistic practice can be justified on grounds that such a practice serves important interests (assuming that this practice can also withstand theoretical scrutiny).

But some arguments advanced by the contextualist imply that having multiple standards of knowledge actually does serve important interests. Consider a case presented by Stewart Cohen:

> Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes I know it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint..." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent. (1999, p. 58)

According to Cohen, Smith's knowledge claim is true when evaluated in his own context (with its weaker standards); but Mary and John are also correct in saying (relative to their context, and the higher knowledge-standard it produces) that Smith does not know.

We noted that the first step in justifying a particular linguistic practice is to show that it serves a legitimate interest. We can see Cohen as arguing that it is important for us to have flexible, contextually-determined standards of knowledge, because in various contexts, the need for a higher (or lower) standard will arise (which need will be determined by our various legitimate interests and needs). This looks more promising than the ordinary language argument for contextualism, which (as we saw) cannot establish the legitimacy of variable standards.

However, it is not enough for the contextualist to argue that having a contextually-determined standard of knowledge serves our interests; and we noted above that we have no quarrel with knowledge_M as exemplified in Cohen's airport case. She must argue that we have interests that are served by making the requirements for knowledge as stringent as is demanded for knowledge_H. For recall the course of our argument: the mere fact that knows_H is enshrined in our ordinary practice does not justify continued use of this notion; we must argue (at the very least) that the use of the content 'knows_H' serves some interest (and survives theoretical scrutiny). What interest can we have in setting the standard of knowledge as high as knowledge_H demands?

One possible way to justify such a periodic ratcheting-up of our epistemic standards is suggested by C.I. Lewis' infallibilist motto: "Nothing is probable unless something is certain." Isn't the need for certainty a sufficient reason to move to an infallibilist standard of knowledge as one's epistemic standard? We have already argued that having and using the term 'knows_L' serves an interest in fruitful communication and distinguishing the epistemic accomplishments of fallible human knowers. But we also have an interest in getting things right, perhaps such that we cannot be wrong–that is to say, we also might have an interest in certainty. It seems that knows_L and knows_H might speak to different interests, where the latter addresses our interest in certainty. The lure of contextualism is that it argues that we can accommodate both sets of interests and intuitions about the nature of knowledge within our epistemic theory.

However, there are at least two, related reasons why this argument fails to justify the adoption of knows_H, even as a complement to rather than merely a replacement for knows_L. First, as we saw above, it is far from clear that any genuine interest is served by a move, even temporary, to knows_H. Adopting knowledge_H as an assumed background for discourse seems to undermine the possibility of making most (perhaps all) assertions whatsoever, taking with it the possibility of communication and meaningful action. How, then, is it interest-serving to retain this particular standard?

But matters are even worse. Let us approach the problem through the issue of *radical* doubt. We will call

doubts radical when their impact is both global and evidence-transcendent *de jure*. To say that they are 'global' is to say that they cast doubt on whole classes of beliefs (maybe all), not just individual ones based on their content or available evidence; brain-in-vat objections undercut all ordinary perceptual beliefs, not just my current belief that the soda can before me is red. To say that they are evidence-transcendent *de jure* is to say that, by stipulation, no evidence that could be available to us could refute the hypothesis raised in the doubt. Evil demon and brain-in-vat arguments are thus ways of raising radical doubts.

The contextualist view is that in some contexts, we properly employ knowledge$_H$, and then in other contexts we properly employ knowledge$_L$. But the idea that context, and our corresponding epistemic entitlements, can shift so easily, fails to engage the radical skeptic seriously. For radical doubt is *explosive*–it cannot be confined to a particular context, and impacts all types of evidence and all types of knowledge-attributions.[14] After considering the evil demon case, the skeptic will charge that we cannot merely retreat to another context and apply another standard of knowledge–for how can we simply 'set aside' these doubts? Certainly, we *do* set them aside, but appeals to ordinary practice don't secure entitlement; that has been a major theme of our argument up until now. Radical doubt cannot be confined to a single context; it bursts out from that context to infect all discourse. If a change in context changes the content of our concept of knowledge, so that in one context we are using 'knows$_H$' and in another we are using 'knows$_L$', the skeptic will want to know what allows us to claim that we have exited the skeptical context. Nor can one set radical doubt aside in ordinary contexts on the grounds that these doubts seem "highly unlikely" or otherwise irrational to believe. Such a position presumes access to prior probabilities that are themselves subject to these doubts. Once we have entered the realm of radical doubt, setting these doubts aside cannot be rational if done for those reasons, even if it is commonly done.

Thus, while we might concede some contextual variability to the standards for proper knowledge-ascription, it is difficult to see what interests might raise these standards all the way to those for knowledge$_H$. We have no communicative or epistemic interest in doing so, and therefore the argument from interest cannot show that

[14] We noted toward the end of section 2 that others (Kornblith 2000, Sosa 2000, Bach 2005, Feldman 2001), argue that contextualism fails to engage the skeptic, because what is at issue between the skeptic and the low-standards purported knower isn't a question of what standard of knowledge is at play, but how good our evidence is (and whether what is present in low-standards cases should count as *evidence* in the first place). The argument we are presenting here is reminiscent of theirs, but importantly different. We aren't conceding the skeptic's argument that we never satisfy the standards of knowledge$_L$; we are arguing that once we admit the legitimacy of radical doubts, once we admit knowledge$_H$ as a legitimate type of knowledge attribution, we cannot simply shift standards to knowledge$_L$ and ignore the skeptical context. It is *this* move, we will argue, that fails to engage with the skeptical challenge.

there is any context in which the notion of an interest can justify raising the standards of knowledge all the way up to those of knowledge$_H$.

To conclude this section, we have seen that the move to include knows$_H$ in the contextualist resolution isn't a happy reconciliation, but an unsound appeal to our interests and a refusal to fully consider the demands of knowledge$_H$ and the explosiveness of skepticism. In the end, we must decide between knowledge$_L$ and knowledge$_H$. And given the necessity of such a choice, and given that 'knows$_L$' is necessary for the very possibility of communication and is the only type of knowledge attribution that allows us to make meaningful knowledge distinctions, 'knows$_L$' is the type of knowledge attribution we should maintain in the language.

5. Fallibilism and Radical Doubts

Thus far, we have argued that there is a pervasive flaw in contemporary contextualist accounts of knowledge: despite their variety, they uniformly make dubious appeals to ordinary practice. But we agree with most contemporary contextualists, however, that whatever the shortcomings of skepticism, it should not be dismissed out of hand. Call this requirement R:

> R: We must explain *both* the plausibility of everyday knowledge ascriptions *and* the relevance of radical doubts with our account of knowledge.

The contextualist response to this is familiar by now: both radical doubts and affirmative everyday ascriptions have their place, and those who raise radical doubts succeed in raising standards in those contexts. We have argued that such a response won't do, and we must develop a more substantial response that concedes the potency of radical doubts while repudiating them in the end.

Note that challenges to everyday knowledge ascriptions often have an escalating character when we do epistemology. S claims that she knows there is a zebra before her; the skeptic suggests it might be a painted mule. This challenge calls into question everyday sorts of perception and recognition. S may then shore up her claim by gathering and/or offering further evidence, such as a closer inspection. The skeptic will typically escalate the attack by calling into question the processes of perception and belief-formation themselves. At some point, we escalate further and move on to radical doubts.[15] Not all general epistemic problems turn on these sorts of radical doubts. The

[15] Some skeptics might reject at this 'escalating' account. They might say that the actual lesson to draw is that we never have any reason whatsoever to accept any proposition, much less know it. Kornblith (2000) has called this "full-blooded skepticism," though it isn't ultimately his view. If so, nothing about our response would change.

Lottery Problem doesn't require any properties that would transcend or elude our efforts to affirm them; rather it suggests that not all our plausible knowledge claims can be compatible with one another, at least not simultaneously. The emergence of radical doubt however is one with a longer lineage, and this strikes us as the most reasonable focal point in understanding R and developing an account that satisfies it.

Earlier, we argued that contextualist accounts appealed too readily to ordinary language and everyday practice. What is the alternative? What we seek in epistemology and in philosophy more generally is an open-ended, diachronic, discursive process in which everyday practices and the interests that drive them shape our theories, but are also subject to rigorous challenge and revision given the right sorts of arguments. Skepticism and radical doubt remain points of philosophical concern because there are enduring tensions in 'common-sense' intuitions about knowledge. No epistemological theory wholly accommodates those intuitions and none could, for they aren't ultimately compatible; the role of epistemology is in part to *force* confrontation with these incompatibilities, both by those within philosophical circles and those outside of them, and make a case for some preferred resolution. Acknowledging this confrontation and seeking a resolution that does optimal (if imperfect) service to our epistemic commitments should not grant a priority to everyday practice that would commit the naturalistic fallacy. Nor should it grant the epistemologist a position entirely outside of our practices from which to legislate how they should be conducted.

Advocates of Cartesian radical doubt will insist that if we accept radical doubts at all, we implicitly accept them for any and every context. Here, we face what Pritchard (2001) has called the "problem of epistemic descent." Once raised, how can we set such doubts aside? Those radical doubts aren't intended as academic problems restricted to particular sorts of discourse. The significance is purportedly context-free and calls on apparently plausible hypotheses about everyday experience and our own imagination. Conceding the relevance of radical doubts without blunting their generality does us no good, and no compelling case that we *should* ignore radical doubts in everyday contexts has been forthcoming. We surely do, but without some plausible set of restrictions, our doing so suggests forgetfulness rather than entitlement. If we are to concede the relevance of radical doubts without simply conceding to the skeptic, we must take a step back to consider what gives such doubts their force. A dynamic between two general interests in ascribing knowledge can explain this.

First, to ascribe propositional knowledge and justification to an agent is to express some resolution on matters that guide and coordinate action. To know that P – whether ascribing it to oneself or to others – is to take

matters as sufficiently settled in favor of P and to assert some authority on how to proceed on that basis. In this sense, knowledge and justification are essentially social affairs. If one knows that P, then whatever secures one's belief can license others to take it up as well.[16] As a community, we may proceed with beliefs that we would not yet call knowledge when we have no other choice, but we would never prefer to do so and to act without such ascriptions is to act with surrender and hedging acceptance, rather than endorsement ("We don't know, but let's try it and see"). Perhaps the best case that can be made for the intuitions many ordinary language philosophers had is that even if the skeptics are right, these dimensions of our shared practices seem indispensable. We could drop knowledge-talk in favor of some preferred-justified-belief status, but privileging some claims or beliefs in these ways is an enduring feature of our shared engagement with the world. (We take this to be Malcolm's argument above.)

Second, agents and communities will generally prefer beliefs that are very likely to be true, perhaps even guaranteed to be true. Note that this was not presumed in describing the previous interest. An agent or community might privilege a claim that is false or less justified than an alternative in coordinating its actions, e.g. many engineers prefer Newton's laws for their simplicity and adequacy for most purposes. Privileging some beliefs short of certainty threatens to leave us caught out, however, with costs that range from mild embarrassment to upending theories and social institutions. Pressure for greater assurance, culminating in a call for certainty, is a common trajectory for the articulation of this demand to take. Dewey noted this theme in much of epistemology:

> Such considerations point to the conclusion that the ultimate ground of the quest for cognitive certainty is the need for security in the results of action. Men readily persuade themselves that they are devoted to intellectual certainty for its own sake. Actually they want it because of its bearing on safeguarding what they desire and esteem. The need for protection and prosperity in action created the need for warranting the validity of intellectual beliefs. (Dewey 1929/1981, p. 381)

Focusing squarely on such efforts to quash all possibility of failure or refutation leads to what McDowell (1995) has called "aggregationist" theories in epistemology, which demand such high levels of justification that they collapse the notion of justification into the notion of truth. Nothing could count as knowledge in such theories unless it came to us in such a secure way that the very idea of refutation would seem nonsensical. If we construe this interest as demanding a *terminal* resolution, not just a *provisional* one, then the skepticism Descartes developed and the foundationalism he proposed to refute it can be seen as of one piece. Indeed, most of western epistemology since the

[16] This may be more or less direct. Others cannot share my perceptions directly, so there are intermediate steps involving my reliability as a reporter, my testimony, etc.

17th century can be seen as driven by some construal of this interest or despair at satisfying it.

These two interests will reach a point of tension. Privileging some beliefs has a practical urgency that we can rarely set aside, and the drive for certainty is one we can rarely satisfy, if ever. The most prominent and explicit consideration of this tension has taken place in debates over fallibilism and infallibilism. Fallibilists have distinguished themselves from skeptics by accepting the general impossibility of certainty, but insisting there were still viable and valuable epistemic distinctions to be made and knowledge to be had. Infallibilists have argued that only beliefs held with certainty could answer the skeptics and resolve other enduring epistemic puzzles. We feel that the fallibilists have the more plausible view, though we will not defend a specific version of fallibilism here. We may address R for any fallibilist theory by showing that the relevance of radical doubt arises from the tension between epistemic interests described above, and that suitable forms of fallibilism may acknowledge that plausibility without undermining knowledge. We claim that the plausibility of radical doubts actually speaks in favor of fallibilist intuitions, rather than against them.

The tension between privileging beliefs and wanting certainty is a persistent feature of our epistemic norms and practices. If an epistemological hypothesis gives us reason to challenge some of our norms and practices and isn't obviously false, then it is relevant to our epistemology and our knowledge ascriptions. Radical doubt precludes almost all forms of certainty, giving us reason to doubt that our intuitions about certainty can be satisfied. Background assumptions supporting radical doubts (e.g. similarities in waking and dream experiences, possibilities of non-veridical experiences) aren't obviously false and therefore radical doubt is relevant to our epistemology and knowledge ascriptions. How should we address it then? We may think of epistemological discourse as the *diachronic* evaluation of our epistemic norms and practices; hypotheses in epistemology become relevant because they bear on that evaluation. This marks a monumental difference between the goals and methods appropriate to fallibilism in contrast with those of infallibilism. Whereas infallibilism would require a *terminal* resolution of questions of justification securing the truth of our beliefs, fallibilism would require a *provisional* one, compatible with our best methods and evidence, but subject to challenge at any time. To an infallibilist, this will sound like an illicit diminution of our epistemic standards, but appropriate articulations of fallibilism will remain rigorously self-critical and self-correcting, and are more squarely grounded in our interests in making knowledge ascriptions.

What to make of the relevance of radical doubt on a fallibilist view isn't a simple matter. But setting aside the skeptic's charges does not have to be dogmatic. It can be a statement that some of the things done with

knowledge ascriptions (i.e. a paramount, if defeasible, authority to assert) are things that we want to keep, even if epistemology shows us that we have to revise our practices to some degree in order to preserve that. To adopt a fallibilist view, we must abandon some features and expectations of our practices and articulate an account that achieves a suitable balance between all our interests in making knowledge ascriptions. (For instance, all practices and beliefs will have to be open and subject to challenge and revision, contrary to much 'common-sense' thinking.) This is no small task, and we should expect any version of it to be a provisional position in an ongoing process. To actually undercut most knowledge ascriptions, radical doubts would have to gain epistemic traction by association with some actual practical difference for us. Proponents of radical doubts might cry foul: if the doubts rest on evidence-transcendent properties and phenomena, then such practical differences are ruled out in principle. But this is false. It seems plausible only if we consider direct evidence, such as our awareness of some property. Indirect evidence might come in cases in which there is some breakdown (i.e. a large number or an apparent global failure of our capacity to anticipate and comport ourselves to the world); there, an inference to radical doubts might be appropriate. Many who experience occasional hallucinations report temporary doubts about the rest of their perceptions, and if such difficulties made our experience persistently irregular, radical doubts would be rationally appropriate. But radical doubts are fatal to our interests under infallibilism. If we hold fast to our interest in reaching certainty, then skepticism clearly prevails, at the costs to our communicative interests we have described. What makes radical doubts philosophically interesting is that they reveal that contradiction between different commitments, and force a confrontation between fallibilistic practices and infallibilistic assumptions. But radical doubts only undermine knowledge in principle if we adhere to infallibilist standards.

Thus, the development of radical doubts as features of epistemological discourse should be seen not as the unfortunate discovery of an inconvenient truth. They constitute an enduring critical response to our infallibilist impulses, our epistemic hubris. Adopting infallibilism has the virtue of restricting our privileged beliefs to those with an unchallengeable measure of authority. If we were in a position to claim certainty and refute radical doubts for critical beliefs, then that would clearly be desirable. But we suspect that this is effectively never the case for beings like us. Even our best epistemic circumstances are ones in which some forms of error are possible, and the only rational response is an ongoing openness to challenge and revision. This extends not only to first-order empirical claims, but also those we reach *a priori* and to the very principles we reach in epistemological discourse. Indeed, because of our fallibility, openness and a set of fallibilist epistemic concepts are far more conducive to

inquiry, and thus serve our interest in truth, better than any infallibilist set ever could.[17]

[17] For more on this connection, see chapters 1 and 7 of our (2016).

References

Bach, K. (2005). The Emperor's New 'Knows.' In G. Preyer and G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning and Truth* (New York: Oxford University Press), 51-89.

Brueckner, A. (2005). Contextualism, Hawthorne's Invariantism and Third-Person Cases. *Philosophical Quarterly*, *35* (219), 315-318.

Cappelen, H. (2011). Against Assertion. In J. Brown, & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 21-48). Oxford: Oxford University Press.

Cohen, S. (1988). How To Be A Fallibilist. *Philosophical Perspectives*, *2*, 91-123.

Cohen, S. (1998). Two Kinds of Skeptical Argument. *Philosophy and Phenomenological Research*, *58* (1), 143-159.

Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Noûs* 33 Suppl (*Philosophical Perspectives 13: Epistemology*), 57-89.

Cohen, S. (2000). Contextualism and Skepticism. *Noûs* 34 Suppl (*Philosophical Issues 10: Skepticism*), 94-107.

DeRose, K. (1995). Solving the Skeptical Problem. *Philosophical Review*, *104* (1), 1-52.

DeRose, K. (2002). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, *52* (4), 913-929.

DeRose, K. (2005). The Ordinary Language Basis for Contextualism, and the New Invariantism. *The Philosophical Quarterly*, *55* (219), 172-198.

Dewey, John. 1981. "Philosophy's Search for the Immutable." In The Philosophy of John Dewey, Vols. 1&2., edited by J. McDermott, 371-388. Chicago: University of Chicago Press.

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

Hawthorne, J. (2004). *Knowledge and Lotteries.* Oxford, UK: Oxford University Press.

Klein, P. (2000). Contextualism and the Real Nature of Academic Skepticism. *Noûs* 34 Suppl (*Philosophical Issues 10: Skepticism*), 108-116.

Kompa, N. (2002). The Context Sensitivity of Knowledge Ascriptions. *Grazier Philosophische Studien*, *64*, 1-18.

Kornblith, H. (2000). The Contextualist Evasion of Epistemology. *Noûs* 34 Suppl (*Philosophical Issues 10: Skepticism*), 24-32.

Lackey, J. (2007). Norms of Assertion. *Nous*, *41* (4), 594-626.

Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, *4* (4), 549-567.

Ludlow, P. (2005). Contextualism and the New Linguistic Turn in Epistemology. In G. Preyer and G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning and Truth* (New York: Oxford University Press), 11-50.

Malcolm, N. (1968). Moore and Ordinary Language. In P. Schlipp (Ed.), *The Philosophy of G.E. Moore, Vol. IV of The Library of Living Philosophers*. Evanston, IL, USA: Northwestern University Press.

McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, *55*, 877-893.

Montminy, M. (2007). Epistemic Contextualism and the Semantics-Pragmatics Distinction. *Synthese*, *155*, 99-125.

Neta, R. (2002). S Knows That P. *Noûs*, *36*, 663-681.

Neta, R. (2003). Skepticism, Contextualism, and Semantic Self-Knowledge. *Philosophy and Phenomenological Research*, *67*, 396-711.

Pritchard, D. (2001). Contextualism, Skepticism, and the Problem of Epistemic Descent. *Dialectica*, *55*, 327-349.

Rysiew, P. (2011). Epistemic Contextualism. The Stanford Encyclopedia of Philosophy. http://plato.stanford.edu/entries/contextualism-epistemology/. Accessed 18 September 2016.

Sellars, W. (1956/1997). *Empiricism and the Philosophy of Mind.* Cambridge, MA: Harvard University Press.

Sosa, E. (2000). Skepticism and Contextualism. *Philosophical Issues*, *10*, 1-18.

Stace, W. T. (1952). *Religion and the Modern Mind.* Philadelphia: J. B. Lippincott Co.

Stanley, J. (2004). On the Linguistic Basis for Contextualism. *Philosophical Studies*, *119*, 119-146.

Stanley, J. (2005). *Knowledge and Practical Inference.* Oxford: Oxford University Press.

Travis, C. (1989). The Uses of Sense. Oxford: Oxford University Press.

Unger, P. (1975). *Ignorance: A Case for Skepticism*. New York: Oxford University Press.

Williamson, T. (2000). *Knowledge and Its Limits*. New York: Oxford University Press.

Wolf, M.P. and J.R. Koons. *The Normative and the Natural*. New York: Palgrave Macmillan.