

Fighting Skepticism with Skepticism:

Supervaluational Epistemology, Semantic Autonomy, and Natural Kind Skepticism

I

The second half of the twentieth century has seen three revolutions in philosophical thought about skepticism. The first was an aspect of the rejection of foundationalism, in both its classical and its logical positivist versions. This broad antifoundationalist movement of thought was powered by the convergence of two traditions: the neo-pragmatism of Quine and Sellars, both heirs to the American pragmatism that culminated in Dewey, and a Wittgensteinian strand followed out by Austin, Bouwsma, and Clarke. These thinkers took it that skepticism arose as a consequence specifically of foundationalist approaches to epistemology and more generally of quite artificial enterprises of high philosophical theory dangerously removed from ordinary practices of assessing knowledge claims, even in the stylized forum of the developed sciences. They concluded that their various critiques of foundationalism and broadly pragmatic theories of meaning either licensed a rejection of the challenge of skepticism or enabled a straightforward response to it.

The 1980's saw a vigorous counter-revolution to this cosy emerging anti-skeptical consensus, spearheaded by Barry Stroud, Tom Nagel, and Bernard Williams. They disputed aetiologies of skeptical worries that trace them to potentially controversial,

highly theoretical philosophical programs such as a quest for certainty inspired by foundationalism. They deny that the skeptic must ignore, replace, or do violence to the ordinary meaning of the term 'knowledge' in order to motivate the distinctive epistemological project of justifying our knowledge wholesale, according to the very highest justificatory standards. They argued persuasively that all the raw materials to which the skeptic needs to appeal are available already in the plain man's ordinary conception of knowledge, even though the epistemological project itself is a highly theoretical philosophical enterprise. Foundationalism is indeed intimately associated with the threat of skepticism, but should be seen as a possible response to that threat, rather than as a basis for it.

The third revolution is the product of the brilliant and original, yet patient and detailed adjudication of the disputes of the first two offered in Michael Williams' book *Unnatural Doubts*. His argument is complex and many-faceted, but it may be thought of as comprising four basic phases. First, he develops and defends a contextualist approach to knowledge. Knowledge claims are always made in a particular context of inquiry, which determines what sorts of claims can serve as or stand in need of reasons, and what counts (in that context) as the sort of conclusive justification needed for claims to know. As Williams deploys it, this epistemic contextualism accords each of the previous contestants an important insight. The antiskeptical camp is correct in pointing out the unique, extreme, and artificial character of the distinctively epistemological context of inquiry within which the skeptic pursues his investigations. The neoskeptics are correct, however, in observing that there is nothing evidently incoherent about the context of

inquiry they construct. And if knowledge claims are in principle context relative, then the mere difference between the epistemological context and more mundane scientific or casual conversational contexts of inquiry does not establish that the subject has been changed, a concept abandoned, or a meaning perverted. “An ordinary concept can impose extraordinary demands when projected into an unusual context.”¹ The epistemological project is the wholesale justification of all of our beliefs about how things actually are, conducted within a context that forbids appeal to any such beliefs in justifying others. Williams further concedes to the skeptics that within this context of inquiry, the conclusive justification of the claims in question—and therefore knowledge—is impossible.

The second phase of Williams’ argument is a diagnosis of the pivotal move on which the skeptic’s argument must rely, as it appears from the point of view of contextualism about knowledge. For what the skeptic can establish, according to the first phase of Williams’ account, is that within a certain context of inquiry, claims of a certain sort can never amount to knowledge. But this does not yet establish that in another context those same claims might so qualify. The skeptic can claim to show that within a certain context, knowledge is impossible. But what he wants is to show, within a certain context, that knowledge is impossible *tout court*. The question is how one can justify the move from

1. Discovering that in context C, knowledge is impossible, to
2. Discovering (in context C), that knowledge is impossible.

¹ *Unnatural Doubts* [Princeton University Press, 1996], hereafter *UD*, p. 166.

This shift of scope is not one to which logic alone will entitle the skeptic. It requires an independent argument privileging the epistemological context of inquiry somehow, so that if knowledge is impossible in *that* context, then it is impossible in *any* context. After all, the fact that the standards of mathematical proof (and therefore mathematical knowledge) are seldom if ever satisfied even in the context of inquiry characteristic of the best behaved empirical sciences, never mind in casual conversational contexts, is not by itself taken to show that knowledge is not to be achieved in those contexts. “A failure of knowledge in one context need not imply a failure in another, still less in all.”² What gives the epistemological context precedence or authority over our assessments of what can be achieved in other contexts?

One of Williams’ central thoughts is that the place on which to focus critical attention in the skeptic’s argument is not on warranting the coherence and legitimacy of a distinctively epistemological context of inquiry, nor on the subsequent establishment of the impossibility of knowledge as it is construed within that context of inquiry, but on the transition from that claim to claims about whether anything properly qualifies as knowledge in any other contexts of inquiry. Sometimes defenders of skepticism simply slide from the one claim to the other. But not all do, and Williams is concerned to evaluate the arguments that are available for the legitimacy of the move. The third phase of his argument is that what stands behind the inference, indeed the *only* basis it can have, is a view Williams calls “epistemological realism.” This is the view first that the kinds of knowledge investigated by the traditional epistemologist—paradigmatically, knowledge of the external world, and directly experiential knowledge—are natural kinds in a sense

² UD, p. 166.

that makes them fit topics for theoretical investigation. The second element of the view is the claim that what picks them out as such kinds is *generic* evidential relations that obtain between them—paradigmatically that all the evidence there is for the whole class of claims about the external world consists of directly experiential knowledge.

The final movement of Williams' argument is then directed at showing that epistemological realism is false. If it is, and if that is the only way the skeptic can move from the impossibility of knowledge within the epistemological context of inquiry to the impossibility of knowledge generally, then the threat of skepticism will have been removed. Further, epistemological realism is recognizably a general form of foundationalism. So if the skeptic's crucial scope shifting move can only be justified by epistemological realism, the claim by the original antiskeptics that skepticism is a consequence of objectionable foundationalist assumptions will have been vindicated. Williams intends to show that traditional epistemology leads to skepticism just insofar as it builds controversial and ultimately untenable foundationalist assumptions into its identification of the *objects* it studies. Williams is a *natural kind skeptic* about epistemological kinds such as experiential knowledge and knowledge of the external world. There are no such things, he thinks—not because we never know things by experience (for instance, that there is something red in front of me), or never know things about the external world (for instance, that I have two hands), but because those bits of knowledge do not fall into epistemologically stratified *kinds* that can be fit topics for theoretical investigation. Williams uses this sort of methodological skepticism to

undercut arguments for substantive epistemological skepticism. This is what I'm talking about in my title.

II

So Williams offers a powerful recasting of the arguments of the New Sceptics, in three stages. First, a sophisticated contextualist understanding of knowledge, together with a defense of the legitimacy of a distinctively epistemological context of inquiry, allows rejection of one central strand of criticism of the notion of knowledge with which the skeptic works. Second, the reconstructed skeptic acknowledges the need to justify a shift of scope, from showing that knowledge is impossible in the epistemological context of inquiry to showing in the epistemological context of inquiry that knowledge is impossible. Third, this shift of scope is justified by appeal to some version of epistemological realism: the view that bits of knowledge come in natural kinds (paradigmatically, directly experiential knowledge and knowledge of the external world) that stand in essential relations of epistemic priority. Sharpened versions of familiar antifoundationalist arguments can then be deployed against the latter notion, to undermine the entire New Sceptical line of thought.

This is a very illuminating way of presenting the structure and lessons of this recent movement of philosophical thought. And I find Williams' characterization and rebuttal of epistemological realism clear, cogent, and compelling. There is one

dimension along which I think his argument may need to be broadened, however. For once the argumentative situation is presented in this new, enlightening way, it seems to me that possibilities open up for arguments that the skeptics have not explicitly exploited—and which, for that reason, Williams has not explicitly addressed. In particular, contextualism about knowledge—perhaps in the end the deepest and most general contribution to epistemology made in this book—provides the resources for a sort of defense of the scope shifting move different from those Williams considers and refutes. Now I think that this other potential route to skepticism is also wrong-headed and can be seen to be defective, and for reasons contiguous with those Williams advances. Seeing what those reasons are and how they bear on the conjectured skeptical strategy requires and enables a generalization of his argument against epistemological realism. Doing that also requires us to be somewhat less irenic and concessive toward the New Skeptics than Williams—largely, I think, for rhetorical reasons—strives to be. For generalizing the diagnosis of the master mistake structuring the view of knowledge that leads to skepticism brings into question the coherence and ultimate intelligibility of the special ^sepistemological^s context of inquiry, which Williams is for the most part willing to concede as legitimate.³

For Williams and the skeptics he considers, the context of inquiry he calls ‘epistemological’ is just one among many—unusual, and even extreme, in many of the demands it makes, but still at the same level as all the rest. When it is so understood, a question naturally arises about why it should be accorded the privilege of determining

³ I use superscripted ‘s’s to indicate ^sscare^s quotes. I explain how I think the semantics of this crucial expressive device works (and so just what I take myself to be doing in using them here) at pp. 545-7 and

what counts as knowledge as such. It is at this point that the need for particular justifications of the scope shifting move arise. But once Williams' penetrating analysis has made this situation clear, another option becomes visible. One might instead seek to build the privileged position of the 'epistemological' context of inquiry (ECI) into its definition, by seeing it as a context at a different level from the rest. What I have in mind is that one might adopt a *supervaluational* account of the 'epistemological' context of inquiry. The claim would be that what we mean by 'knowledge', strictly speaking, is just whatever would count as knowledge from the point of view of *any* and *every* context of inquiry. Such a view might derive its historical philosophical credentials as a reading of Descartes' epistemology-initiating resolve to doubt everything that can be doubted. The thought would be that if for some claim p there is any point of view (context of inquiry) from which p is dubitable in a sense that precludes its counting from that point of view (or in that context of inquiry) as knowledge, then it is dubitable full stop, and precluded from counting as knowledge, full stop. The supervaluational senses of 'dubitable' and (so) 'knowledge' result from quantifying over (first order) contexts of inquiry. In this supervaluational sense of hyperbolic doubt, one need not find some one method (first order context of inquiry) that permits one to doubt at once all claims of a certain class (find that their epistemic credentials fail to establish their truth). It is enough if for each claim there is some context of inquiry in which it can be doubted. What is knowledge "strictly speaking" is what there is sufficient evidence for by the strictest standards of evidence; but there need not be some one set of strictest standards of evidence, one first order context of inquiry within which all the evidence for all the claims is to be assessed. Rather, by "strictest standards" we mean something intelligible only by quantifying over

588-90 of *Making It Explicit* [Harvard University Press, 1994].

all such ordinary contexts of inquiry: if *any* context of inquiry has standards of evidence so strict *with respect to this claim* that it cannot be established there, then by the strictest standards applicable to it, it cannot be established.

Such a supervenient understanding of the ECI precludes the need for a separate justification of what on other construals appears as the scope shifting move. For the *supervenient epistemological context* (SECI) is defined as the second order context in which something counts as knowledge (as sufficiently warranted) just in case it counts as knowledge in *every* first order context. Once the possibility of this way of proceeding has been broached, the skeptic might appeal to it as capturing the real intent of a move that Williams considers, and rightly rejects, for justifying the scope shifting move. For one attempt at justifying that move appealed to the ECI as privileged in being “purely theoretical”, in that it abstracts from all practical concerns, including those derived from the economics of inquiry (such as Simon’s satisficing). The problem with this proposed defense, Williams points out, is not that it does not justify a distinctive context of inquiry—one, indeed, that deserves a certain sort of privilege or priority with respect to the others. It is rather that it cannot be identified with the ECI as required by the skeptic. For the ECI excludes far more than merely practical concerns. In the ECI one is not allowed to appeal to any beliefs or claims, including theoretical ones, as evidence in establishing knowledge claims, if it is of the same kind as the one whose credentials are being established. And being of the same kind in the relevant sense means being capable of being brought into question by the same wholesale method of doubt (e.g. appeals to the possibility of dreaming, or of the subject’s being a brain in a vat). It

is easy to see both that what is being ruled out as evidential extends far beyond the merely practical, and that the diagnosis of reliance on epistemological realism is apt. But having seen this telling response, the skeptic may be tempted to recast the initial argument in the context of a supervenational construal of the ECI. The thought would be that for various practical reasons peculiar to itself, each first order context of inquiry sets different evidential standards, lax along some dimensions, more strict along others. (There is no need for commitment to there being only one dimension of strictness.) It is *these* practical considerations that the SECI abstracts from, and in *this* sense that its understanding of the concept of sufficient evidence is “purely theoretical.”

The two points worthy of notice are that the supervenational understanding of the epistemological context of inquiry achieves the effect of the scope shifting move, and that it does so without even implicit appeal to epistemological realism. The SECI is not just one context of inquiry among others. It defines a notion of knowledge, strictly so-called that is superordinate to and derived from those defined with respect to first order (‘ordinary’) contexts of inquiry. Its notion of strictness is made intelligible, not by appeal to some peculiar and hard to motivate project, but just in terms of what is already made available by the ordinary contexts of inquiry. Its restrictions on what counts as sufficient evidence do not depend on there being epistemological *kinds* of claims, all of which must be justified or put in question together. It requires no commitment to epistemological realism, and hence none to the sort of foundationalism Williams rightly sees as implicit in views that do involve such commitment.

Further, given some not unreasonable assumptions, the supervenient way of defining a strict sense of ‘knowledge’ will underwrite a thorough-going skepticism. For with the possible exception of what Wittgenstein (in *On Certainty*) calls “hinge propositions”⁴—that is, for a very large class of claims indeed—it is plausible that for every p there is some context of inquiry C such that p cannot be conclusively established in C . If so, then in the supervenient sense, p cannot be *strictly speaking* conclusively established, and so, strictly speaking, not known. We do need to be a little careful here. If contexts of inquiry are allowed to be defined narrowly, so as only so much as to address a restricted range of concerns—perhaps because of restricted practical interests—then on this formulation skepticism would be bought so cheaply as to be of no philosophical interest, and would lose all connection with traditional epistemological projects. For if the metallurgy lab counts as establishing a context of inquiry, its lack of evidential resources for addressing the question of how many legs horses typically have would entail that we cannot, strictly speaking, know the answer to that question. Such a notion of knowledge, strictly so-called, is of no interest whatsoever. If the concept context of inquiry is to be understood so liberally, then the superveniently strict sense of ‘sufficient evidence’ should be reformulated, so that one knows in the strict sense only what can be known in the ordinary sense in every context of inquiry *that addresses that issue*. (Williams never tries to delineate very precisely what context of inquiry is, possibly as a result of his aim to bring order to the discussions of such a variety of

⁴ Williams himself seems—correctly, in my view—to treat the possession of such ‘hinge’ status by propositions as context relative.

philosophers writing on different epistemological issues and from so many different perspectives.)⁵

It appears, then, that thinking about knowledge in the contextualist terms that Williams introduces opens up the possibility of a route to skeptical conclusions that does not depend on epistemological realism. If so, Williams has refuted one version of skepticism, only to make possible a novel, mutant strain, immune to the conceptual antibiotics that sufficed to eliminate its more familiar variants. In fact I think the supervenient conception of knowledge in the strict sense is incoherent and unsustainable, and hence no threat to Williams' conclusions. Seeing why and how it leads to a generalization of the considerations he advances in support of those conclusions.

III

We might begin by considering another application of the supervenient machinery: to lay alongside notions of knowledge, sufficient warrant, and conclusive evidence in a *strict* sense, dual complementary conceptions of them in a

⁵ This point might usefully be generalized, by introducing a parameter of *admissibility* of a context. The *relativized* strict supervenient approach to knowledge would count a true claim as expressing knowledge if and only if it is conclusively warranted in every admissible context, and as not knowledge iff there is some admissible context in which it is not conclusively warranted. It is then a short step to relativize this parameter itself to contexts, talking about what other contexts are admissible with respect to or from the point of view of a given one. The result would be to introduce an *accessibility* relation among contexts of inquiry. Knowledge claims in the strict sense in a context would then correspond to *necessary-knowledge* claims in that context: claims that count as known in every context accessible from the index context. The set of (first order) contexts of inquiry plus the accessibility relation among them would then underwrite a semantics for a modal logic of knowledge claims of the familiar Kripkean sort. Depending on the philosophical motivation and characterization of the significance of the accessibility relation, one might

supervaluationally *lax* sense. In this sense, one would count as knowing that p if there is *any* first order context of inquiry in which one counts as knowing that p (having sufficient warrant or conclusive evidence for it). Now the skeptical epistemologist can readily argue that these lax notions are not of interest in the same way or for the same reasons that the corresponding strict notions are. I actually think that a debate on this issue would be philosophically revealing in various ways, but I am not going to pursue that issue. My concern here is rather with the likelihood that for many claims, in this lax sense we could be said to know both that p and that $\sim p$ —e.g., “Electrons are particles.” It is hard to argue for this in a careful way, since we don’t have a very definite notion of context of inquiry to work with. But since contexts are to differ precisely in the sorts of considerations they allow to count as evidence for and against claims, it is not easy to see how the possibility could be ruled out that one context admits a set of endorsed considerations as evidence sufficient to establish p , while not considering or giving weight to others, which in some other context would count as sufficient (even perhaps when conjoined to the first set of considerations) to establish $\sim p$. Depending on what motivates interest in the lax senses of ‘know’ and ‘conclusive evidence’, the fact that we could know or have conclusive evidence in the lax sense for both p and $\sim p$ might not be a cause for concern. After all, one of the automatic benefits of the supervaluational construction is that it ensures that so long as no first order context of inquiry certifies the contradiction $p \& \sim p$, neither will the lax disjunctive supervaluation. But the fact that one might worry on structural grounds about whether the disjunctive supervaluational sense

then debate the relative merits of taking T or S4, say, as the relevant logic. I will not discuss these generalizations of the supervaluational approach to philosophical epistemology introduced here.

of ‘know’ applies to too *much* raises the complementary question as to whether the conjunctive supervenational sense applies to too *little*.

Too little for what? Too little to be a *semantically autonomous* context of inquiry—because (as we will see) it is too little to be capable of fixing the meaning or content of the claims or beliefs being considered as candidates for knowledge. The strict supervenational context meets a version of the skeptical epistemologist’s condition that no collateral beliefs be available to appeal to as evidence for a target belief if it is possible to raise the same sort of questions about their justification as it is about the justification of the target belief. I am claiming, however, that it does so without resting on a substantive (and so epistemologically realistic) conception of “same kind”; the meaning of that phrase is fixed supervenationally, in terms of a sufficient challenge in some first order context or other. Because the strict supervenational context meets this condition, it meets the condition that Williams rightly summarizes like this:

When, in an attempt to philosophize, we are required to prescind from all situational factors, we are, in effect, asked to specify a basis for asserting a proposition given nothing beyond the semantic content of the proposition itself.⁶

It should be no surprise that the requirement that sentences be justifiable by appeal to their meaning alone, without appeal to the truth of any other sentences, then yields a skeptical result. I’ve suggested the supervenational context as a way of making this demand intelligible, without appeal to epistemological kinds.

But this strategy has a cost, too. For I think it can be shown (though I will not attempt to do so here) on relatively familiar Quinean grounds that the *meanings* of ordinary empirical propositions are not so much as intelligible apart from consideration of what (else) is *true*. If so, then prescindng from all commitments collateral to a given belief is removing an aspect of the context that is essential to the belief's having the content it does—or indeed, any empirical content at all. The claim is not that for any content there is some set of collateral beliefs such that one cannot entertain the content unless one has those very beliefs. It is rather that for any content, one cannot entertain it unless one has some set of collateral beliefs or other. Could we mean lion (or anything intelligible in the vicinity) by any of our words if it is not true that there have been mammals, that some physical objects exist, that it is possible sometimes to see neighboring middle-sized objects...? Perhaps for each such collateral belief one can tell a story, but not one story for all possible such auxiliary hypotheses. This is a point that the supervaluational machinery embodies. For in *each* first order context in which the epistemic status of the claim, say, “People have on occasion seen lions,” is evaluated, there are *many* collateral beliefs with respect to which its credentials are assessed and its meaning articulated and stabilized. If the strict supervaluational context really represented everything that is true in *every* first order context, it would also include such a set of collateral beliefs.⁷ But since which other beliefs form the semantic background of a given believable content may (we are supposing) vary from context to context, the

⁶ *UD*, p. 171.

⁷ Or truths, depending on how we think about such contexts. The point is that in every first order context in which *p* is known, there is a set of collateral beliefs that are not *strictly* known, to which it is evidentially related. That is not true in the supervaluational context of inquiry.

supervaluational machinery transmits none of them to the strict supervaluational context resulting from their intersection.

Indeed, it is not clear that it makes sense to talk about a supervaluational *context* of inquiry at all. For contexts of inquiry, as Williams deploys that concept, do not consist just of sets of *sentences*, propositions, or claims. If they did, then the intersection (or union) of any such set would itself be such a set, and so a candidate context. Contexts of inquiry include also concrete skills, practices, and procedures, which help give sense to the verbal elements of the context. Individuating these across changes in context will be a challenge for the would-be supervaluational skeptic. Further challenges involve showing that there are any such background practices in common to all the other contexts of inquiry in question, and again that, if so, what is common is sufficient to make the verbal claims concretely meaningful.

There are three ways the strict supervaluationist might think about the contents of the propositions she considers, to resist this charge of semantic parasitism. First, she could be resolutely semantically atomistic, taking it that the contents of each claim are totally independent of the endorsement of any other contentful claims. Second, she could take it that the residual least common denominator of the contents of a particular claim in all the ordinary contexts itself suffices to be a sort of content—one that should be recognized as the core of the meaning in all those contexts. Third, she could claim that the meaning of ordinary empirical concepts can be elaborated purely in terms of how things *appear*, without commitment to how anything actually *is*. I think none of these

semantic approaches, neither the atomist, the LCD, nor the phenomenalist strategies, is finally workable, and I would appeal to the considerations Sellars advances in “Empiricism and the Philosophy of Mind”⁸ to argue that. But I won’t try to do so here.

It is because the epistemological context rules out appeal to *any* background beliefs—the totalizing condition, which insists that in this context the epistemic credentials of all our beliefs be assessed together, at once, wholesale—that a skeptical conclusion results. But if it can be shown that leaving out all the background beliefs removes something *semantically* essential to the meaning or content of the very beliefs whose credentials are to be assessed, then the *epistemological* significance of this consequence becomes questionable. For it will then emerge that the strict supervenient context is semantically parasitic on the first order contexts of which it is a product. It is only because of their role in contexts in which ordinary knowledge claims are endorsed that concepts have the contents they do—or any contents at all. This would not mean, I think, that the strict supervenient version of what Williams calls the “epistemological context” need be convicted of incoherence, or of failing to qualify as a context of inquiry. In the absence of more constraints on the latter notion, it is hard to tell. On this issue, it would give us reason perhaps to endorse Williams’ most careful statement of his for-the-sake-of-argument irenic attitude on this point:

I have not denied that the request for a completely general understanding of human knowledge *makes sense*. Or, as I would prefer to say, I do not deny that we can make a certain amount of sense of it.⁹

⁸ Reprinted with an introduction by Richard Rorty and a study guide by Robert Brandom [Harvard University Press, 1994].

The real question would seem to be not whether the strict supervenient context is an intelligible context, so that its notion of knowledge is a genuine one, but why we should care whether or not certain sorts of knowledge are possible in this context. What difference does or should it make that in this strict sense knowledge of a set of claims is not possible, if the contents of those claims are themselves only intelligible in other contexts, which fund notions of knowledge that do extend to some of those claims? This issue will not be nearly as urgent if the epistemological context can fund *some* propositional contents for empirical claims, even though these may be quite different from the contents we ordinarily take ourselves to be entertaining and endorsing. But if (as I believe) it can fund *nothing* recognizable as a propositional or conceptual content, the case is different. For then the intelligibility of the strict sense of knowledge (which the skeptic withholds from ordinary empirical claims) is itself semantically parasitic on the granting of ordinary knowledge claims.

This is the situation that readers of “Empiricism and the Philosophy of Mind” should expect. For there Sellars argued that a certain kind of epistemological foundationalism is harmless—he was inclined to think, even platitudinous. One can grant the evidential priority and privilege of noninferential reports, as compared to purely theoretical claims (that is, for Sellars, claims that can only be appropriately asserted as the conclusions of inferences), so long as one recognizes that they do not form a semantically autonomous stratum. The capacity to make noninferential reports, even of how things merely appear to one, depends upon and is intelligible only in a context and

⁹ UD, p. 223.

against a background that includes specifically inferential practices: practices of endorsing claims that do not simply report what one noninferentially perceives. If that is right, then there is no intelligible position in which we know only what we directly perceive or experience, and have no basis at all for moving inferentially to claims that go beyond that basis. From this point of view, what is pernicious about epistemological kinds is unsustainable assumptions about their semantic autonomy. There is nothing wrong with acknowledging that token beliefs acquired noninferentially or statements of how things merely appear to us share certain epistemically significant properties, so long as we do not assume that they are intelligible as contentful apart from all relation or commitment to inferentially acquired beliefs or what does not merely appear. The soft underbelly of skeptical epistemology is its implicit semantics.

IV

My discussion so far depends on Williams' reframing of the debate between contemporary skeptics and anti-skeptics against the background of a contextualist approach to knowledge. I have suggested that once things are seen in this new light, space opens up for a new skeptical strategy. And I have sketched a response to that strategy that is different from the one Williams offers to all the skeptical arguments he considers—arguments which have actually been defended by sophisticated contemporary

skeptics. I don't see this response merely as supplementing Williams' diagnosis and critique of epistemological realism, however, but as potentially deepening our understanding of it. For I think that the most powerful arguments against epistemological realism proceed by unpacking its dependence on assumptions of semantic autonomy.

We can see this in two stages. In broadest terms, the issue Williams raises is the suitability of such generic categories as experiential knowledge, knowledge of the external world, and knowledge as a whole to serve as objects of systematic theory. The motivation for thinking of experiential knowledge and knowledge of the external world as complementary kinds, partitioning the field of knowledge, stems from *separability* arguments, what amount to arguments for the *independent variability* of the two components. Classical skeptical global extensions of the argument from illusion, such as dreaming and brain-in-a-vat thought experiments, seek to show that we can hold one component (^sexperience^s) fixed, while wildly varying the other (the external cause of sense and goal of intellect). But the idea that the least common denominator in such cases suffices to determine an intelligible notion of *content*—the idea that there is a kind of content that deserves to be called 'narrow', by contrast to the 'wide', world-involving sort—is not semantically obvious, innocent, nor in the end, I think, defensible. Substantial assumptions about semantic autonomy are built into the supposed discrimination of *experiential knowledge* in this sense as a kind of knowledge (or belief), and hence into the discrimination-by-contrast of the complementary kind of *knowledge of the external world*.

The second point has to do with the idea of addressing the issue of justification of our knowledge in a wholesale fashion: thinking about how to justify knowledge-as-a-whole. (The previous point involved subdividing it into epistemological kinds, one of which needed special justification, the other of which did not.) Here Williams' contextualism comes to the fore. If what is required for an adequate or conclusive justification of a claim varies from epistemic context to epistemic context, it cannot simply be taken for granted that there is some common property of being justified or relation of justifying that they all exemplify, and which is a fit subject for systematic investigation. Of course, context relativity as such does not rule out being a suitable object of theory. Indexicality, for instance, is a kind of context relativity that admits of theoretical codification. Williams' *deflationism* about the concept of knowledge turns, I think, on the particular nature of the context relativity of knowledge claims.

With an indexical expression such as 'here', one needs to settle the index in order to know *which* place is indicated, but not in order to know what a place *is*. Places can in principle be identified and individuated without appeal to the indices to which the reference of tokens of the type <here> are relativized. But in the case of the context relativity Williams claims for knowledge, the relevant context is a set of beliefs that help determine the *meanings* of the claims whose credentials are to be evaluated. Apart from some such context, it makes no sense to talk about what is supposed to be evaluated in that context. There are no *autonomous semantic units*, which can be taken to be justified under one set of circumstances in one context, in under another in a different context. The sentences are repeatable across contexts. But they are not what we know (except in a

secondary sense). They are only expressions that tag the propositional contents that are the primary candidates for the status of knowledge. If the attribution of content to any belief only makes sense in a context that includes commitment to many other such contents (though not always the same ones), then there is no context in which commitment to all our beliefs could simultaneously be assessed. Knowledge as such is not a fit topic for investigation—for fundamentally *semantic* reasons.

V

Williams is a deflationist about *knowledge* in much the same sense (though not for the same reasons) that others are deflationists about *truth*. His skepticism about the objects of epistemological inquiry pre-empts the sort of skepticism that purports to result from such inquiry. He is skeptical about such categories as experiential knowledge, knowledge of the external world, knowledge as such, and knowledge as a whole. He is not in general skeptical about the particular items that are alleged to belong to these types. If ‘experiential knowledge’ means something more than knowledge that is immediate in the sense of being noninferential, I think one *ought* to be skeptical even about the existence of instances of it, and I think Williams agrees. But I do genuinely know (at least sometimes) how things appear to me, even when I don’t know how they actually are. And sometimes I know things that are paradigms of the complementary category of knowledge of the external world, such as that I have two hands and that there

have been black dogs. Williams is skeptical about *knowledge*, but he doesn't deny that we know lots of things.

I think it is helpful in understanding this sort of *natural kind skepticism* to consider some other examples. Think, for instance, about the corresponding attitude as applied to *morality*. One standard way of understanding the central challenge for moral theory is as saying what is distinctive about moral *reasons*—to explain the nature and origins of their bindingness, of the sort of commitments and entitlements they articulate, in relation to other sorts of reasons for action (and belief). One might take it, however, that although what are usually denominated 'moral' reasons for action are, at least often, genuinely *reasons* for action, there is nothing in particular that such reasons have in common that would justify grouping them together as *moral* reasons. Features that have been appealed to in an attempt to demarcate them in a principled way—universality, unconditionality, overridingness, and so on—seem to be both too narrow and too wide, and don't permit the specification of individually necessary and jointly sufficient conditions for membership in what turns out to be a motley, gerrymandered concept. One might even suspect that the concept of distinctively moral reasons is a historical relic, an artifact of a philosophical outlook and project that belongs to an age we have rightly moved beyond—in the case of the moral, motivated by the (no doubt laudable, but by now merely quaint) attempt to secure by secular means a successor concept to a notion of a kind of bindingness previously associated with divine commands. An adherent of such a view might seek to understand the concept good reason for action without any

antecedent commitments regarding the existence of a distinctive subclass of *moral* reasons.

In a similar vein, one might be a natural kind skeptic about the *aesthetic*. One need not deny that there are any beautiful objects, nor that some beautiful things are valuable, in order to deny that beauty is a fit topic for systematic inquiry, or that there is any distinctively aesthetic sort of value or feeling. There need be no such thing as art (artists, works of art) as such. Perhaps all that ties together the diverse particulars we classify under these headings is specifiable in terms of their role in various social institutions.

Indeed, it seems to me that the time may well have come to be natural kind skeptics about *all* of the late eighteenth century (Kantian) categories that sometimes seem so unavoidable (and almost invisible) to us: theoretical (epistemological), practical (ethical), and aesthetic. It may be a salutary exercise to train ourselves to do without these categories—creatures as they are of theoretical frameworks we can hardly endorse today—while not ceasing to make the sorts of specific evaluations they seek to systematize. Such a policy would at least have the advantage of forcing us to rethink the tired rhetoric that shapes contemporary debates between neo-enlightenment defenders of reason, knowledge, and the search for hard facts on the one hand, and neo-romantic defenders of feeling, play, and the creation of fashionable novelty, on the other.

Of all the Kantian kinds, knowledge has seemed the best established, and least avoidable. One of the lasting achievements of Williams' path-breaking work is to begin to teach us how we might do without it.

Bob Brandom

University of Pittsburgh