Knowledge Second (for Metaphilosophy)

Jennifer Nado
University of Hong Kong

The concept of knowledge occupies a central role in epistemology – so much is obvious. But knowledge and its compatriots – evidence, justification, warrant, and the myriad of other concepts intimately connected with knowledge – also dominate the comparatively recent literature surrounding the use of intuitions in philosophy, and surrounding the philosophical consequences of experimental philosophy. Indeed, both experimentalist critics of intuition and their opponents have a tendency to state their conclusions using the standard epistemic concepts of mainstream analytic epistemology. This has, quite unfortunately, led to a misleading picture of experimentally-based criticisms of intuition as advocating 'intuition skepticism' – in other words, as denying that intuitions are a source of knowledge. This, in turn, has led to a metaphilosophical literature dominated by debates over whether such a complete rejection of the evidential value of intuition is reasonable, or even possible.

But concepts like knowledge, I'll argue, are in fact ill-suited for purposes of an evaluation of philosophical methodology. This is because the epistemic standards that govern inquiry in philosophy are separate from, and plausibly much more stringent than, the standards one must meet in order to know. This is not intended as a contextualist claim; the claim is not a semantic one, but one concerning the epistemic norms governing philosophical inquiry. The existence of these separate norms, in turn, invites a pluralistic approach to epistemic states. We might, for instance employ counterpart terms – evidence$_P$, justification$_P$, knowledge$_P$, and the like – to refer to the epistemic states that result from meeting the epistemic standards generated by philosophy-

specific norms.[1] We can then note that the fact that a given method or process generates evidence, justification or knowledge does not guarantee that it will generate evidence$_P$, justification$_P$, or knowledge$_P$. A proposition might well be known without being known$_P$; it might be evidence without being evidence$_P$.

The appropriate conclusion to draw from experimental philosophy studies, I'll claim, is best stated in terms of philosophy-specific epistemic categories like these, rather than the categories familiar from standard analytic epistemology; experimental philosophy, then, would benefit from a pluralistic conception of epistemic states. Given such a conception, we can then gloss the experimentalist conclusion as follows: though intuitions may produce knowledge, experimental findings indicate biases and epistemic flaws that suggest that our current use of intuition may not produce knowledge$_P$.

This, I would suggest, is a much less problematic stance than 'intuition skepticism'. The restated experimentalist challenge I offer can be viewed as analogous to the uncontroversial claim that the biases and epistemic flaws that afflict ordinary, uncontrolled observation render it insufficiently reliable for use in an experimental context: thus making epistemically obligatory, within the sciences, the use of e.g. measuring instruments and procedures such as double-blinding. It's obvious that such claims can be made without threatening skepticism; *mutatis mutandis* for philosophy-specific criticisms of the use of intuition.

The reorientation of our metaphilosophical framework that I'll be suggesting implies that experimentalists and defenders of intuition are, in fact, frequently talking past one another; while

[1] In the case of knowledge$_P$, it would presumably also be a requirement that the content of said state is true – though even this might be questioned if one were, e.g., inclined to some form of anti-realist interpretation of philosophical inquiry (which I am not, I might add).

defenders busy themselves with arguments about intuition's ability to justify belief or to generate knowledge, such defenses fail to address the real challenge that experimental findings potentially present to our current methodological practices. Ultimately, all this points to a more general moral: the monolithic focus on 'knowledge' as the primary epistemic state of interest in philosophy obscures many of the subtleties involved in evaluating our multifaceted practices of inquiry, both in philosophy and in other fields. There exist many different epistemic standings of interest; it is, I think, a losing game to attempt to capture all these with any reasonably unified account of knowledge. Worse, it is a distraction from the more central question of how to best investigate reality.

1.

Elizabeth O'Neill and Edouard Machery, in their introduction to a recent volume on experimental philosophy, present the experimentalist argument against intuition as follows:

> "1. The judgments elicited by philosophical thought experiments are significantly influenced by factors that do not track the fact of the matter.
>
> 2. If a judgment is significantly influenced by factors that do not track the fact of the matter, that judgment is not reliable.
>
> 3. If a judgment is not reliable, it cannot provide warrant for assuming its content.
>
> 4. Hence, the judgments elicited by thought experiments do not provide warrant for assuming their content." (Machery and O'Neill 2014, xvi)

The first premise of this argument is, of course, meant to be supported by the various findings generated by experimental studies on intuition – findings that suggest that intuitions may be subject to influence by cultural background (Weinberg et al 2001, Machery et al 2004), emotion (Wheatley and Haidt 2005; Schnall et al 2008), order (Swain et al 2008), personality (Feltz and Cokely 2008, 2009), gender (Buckwalter and Stich 2014) and so on. One could object to this premise on grounds of methodological flaws in the experiments (see e.g. Kauppinen 2007, Cullen 2010, Woolfolk 2013), or by noting that several of the aforementioned studies have failed to replicate (Nagel et al 2013, Adleberg et al 2014, Seyedsayamdost 2015). But set such issues aside for present purposes – though the current body of experimental studies no doubt has its fair share of flaws, it's reasonably plausible that a more mature experimental research programme may one day fully vindicate premise 1.

More important, for our purposes, is the dramatic claim expressed by the conclusion – the epistemic flaws exhibited by thought experiment judgments entail that the judgments provide *no warrant*. This characterization of the experimentalist argument is in no way atypical. Consider the following statements by members of the experimentalist camp:

> "Sensitivity to irrelevant factors undermines intuitions' status as evidence" (Swain, Alexander and Weinberg 2007, 141).
>
> "Experimental philosophy challenges the usefulness of [appealing to intuition] in achieving justified beliefs" (Alexander, Mallon and Weinberg 2010, 298).
>
> "The ultimate hope is that we can use [experimental philosophy findings] to help determine whether the psychological sources of the beliefs undercut the warrant for the beliefs" (Knobe and Nichols 2008, 7).

> "[Susceptibility to bias gives us] good reason to think that intuitions are unreliable" (Tobia, Buckwalter and Stich 2013, 631).
>
> "[Experimental findings are] bad news for the standard philosophical assumption that the contents of people's intuitions are very likely to be true" (Stich and Tobia 2016, 13).

Such statements are quite strong. The first three explicitly deny that intuition produces the standard 'epistemic goods' that epistemologists urge us to pursue – goods such as evidence, justification, and by extension knowledge. The latter two imply a similar denial given the extremely plausible assumption that substantial unreliability is incompatible with the production of said goods.

Opponents of experimental philosophy, too, characterize its claims in similar terms. Max Deutsch claims that "[experimental findings] should lead us, or so one might think, to question whether we are properly justified in intuiting, and then believing, what we do about the cases" (2015, 7). Elijah Chudnoff claims that experimental philosophers believe their results to "give us reason to think that intuitions are unreliable" (2014, 108). Boyd and Nagel (2014) portray experimental philosophers as arguing that "we have empirical evidence for the unreliability of epistemic intuition" (pp. 114). Finally, one has the bevy of defenses of intuition that take, as their primary target, a 'skeptic' about intuition – implying that the intended opponent denies that intuitions produce knowledge. Instances of this are found in, for example, the work of George Bealer (1992, 1996), Joel Pust (2000), David Sosa (2006), Alvin Goldman (2007), and Timothy Williamson (2004, 2007, 2016).

Framing the epistemological consequences of experimental philosophy in terms of these 'standard' epistemic categories – evidence, justification, knowledge, and so forth – is puzzling. It

suggests that what is primarily at issue is intuition's 'overall' epistemic status. If intuition is a source of evidence/justification/knowledge, then the experimentalists lose; if it is not, then they win. But methodological critiques in other fields don't at all take that form. Early proponents of double blinding in science, for instance, were not 'skeptics' about ordinary observation; they did not take their primary question to be whether or not observation generates knowledge. The focus on knowledge and its epistemic siblings is, I'll argue, just as inappropriate for debates over philosophical method. The threat of skepticism generated by said focus is wholly illusory.

Timothy Williamson's complex and powerful critique of 'intuition skepticism' provides, I think, a perfect illustration of why this is so – though the argument succeeds on its own terms, ultimately, I claim, it misses the point. In what follows, I'll use Williamson's critique as an example of how metaphilosophical arguments framed in terms of the concepts of standard analytic epistemology are, quite generally, a red herring. An adequate evaluation of the consequences of empirical findings on intuition, I'll argue, simply cannot be conducted in such terms; finer carvings of epistemic space are needed.

Williamson's critique of intuition skepticism is multi-dimensional, stemming in large part from various doubts he has about the use of 'intuition' as a defining feature of philosophical methodology. One such doubt involves a general anti-psychologism with regard to evidence; on Williamson's view, we should conceive of our evidence in philosophy not as consisting of propositions like 'I intuit that the Gettier case is not a case of knowledge', but as consisting of propositions like 'the Gettier case is not a case of knowledge'.[2] Psychologizing our evidence, Williamson notes, invites skepticism – and indeed, external world skepticism derives much of its

[2] Throughout the paper, I will assume (with Williamson) that evidence consists of propositions.

force from the difficulty of arguing from psychological premises alone to any non-psychological conclusion whatsoever. We should, then, resist the temptation to characterize philosophical evidence in terms of intuition.

Williamson is plausibly correct here, but note that this by itself isn't sufficient to render experimental philosophy problematic - so long as its proponents are relatively careful with how they frame their critique. A scientist's evidence plausibly consists of propositions about e.g. acids, proteins, or what have you; it does not plausibly consist solely of propositions about the perceptions experienced by scientists in the lab. Yet it's obvious that epistemological deficiencies in perception suffice to problematize certain methods in the sciences – hence the use of e.g. measuring devices rather than bare unaided observation. A methodological critique based on experimental evidence for similar deficiencies in intuition in no way requires psychologizing evidence.

A second aspect of Williamson's critique, however, appears *prima facie* to be more problematic for the experimentalist. On Williamson's view, the judgments we make in response to thought experiments are underwritten by multiple different cognitive processes, including ordinary counterfactual reasoning and the quite general ability to apply concepts to particulars. Quite obviously, such abilities also underwrite much of our everyday cognitive activity. As Williamson puts it, "any psychological kind that includes armchair philosophical judgments includes a mass of non-philosophical judgments too" (Williamson 2009, 475). He notes that "the obvious danger is that the category of philosophical intuitions will be stretched so wide, encompassing virtually anything one says about actual cases, that experimental philosophers' critique of reliance on philosophical intuitions will become a global skepticism" (Williamson 2016, 24).

Such observations bring out two separate problems with the current focus, within metaphilosophical debates, on intuition's 'overall' epistemic status – that is, its status as a source of knowledge, justification, and/or evidence. The first is that, though experimentalists aim their critiques at philosophical inquiry, intuition is clearly not limited to that practice alone. 'Intuition' covers states that are used in ordinary cognition, in the sciences, and so forth, and so a wholesale ban on intuition renders those activities problematic as well. If intuition does not generate knowledge (or evidence, or what have you), then a great deal of our everyday judgments are suspect. The second problem is that 'intuition' covers states that are heterogeneous – the single term 'intuition' obscures the fact that philosophical judgment invokes a variety of quite different forms of cognition which quite plausibly differ widely in their epistemological features.[3] To urge the rejection of 'intuition', then, is to urge the rejection of a number of quite different psychological processes in one fell swoop. I would add that even the aforementioned 'ability to apply concepts to particulars' isn't plausibly a single ability, but many. Consider the experimental evidence suggesting that disgust affects moral judgments (see e.g. Wheatley and Haidt 2005; Schnall et al 2008); consider how incredible it would be to argue, on that basis, that we should assume logical judgments will be similarly affected. Yet this is exactly what experimentalists in effect suggest when they make claims like those expressed in the quotes above - they implausibly treat intuition as an epistemologically uniform kind, and they implausibly lump together as 'unreliable' or 'unjustified' a variety of presumably epistemologically dissimilar states.

Ultimately, the heterogeneity of intuition is - like an anti-psychologistic account of evidence – relatively easy for experimentalists to accommodate. Rather than argue that

[3] See also Nado (2014).

'intuition' is unreliable or fails to produce evidence, the experimentalist merely needs to draw her conclusions more narrowly – in terms of 'moral judgment' or 'logical judgment' rather than 'intuition', for instance.[4] This does, I think, give the experimentalist the resources to avoid a blanket global skepticism: rather than condemning "our practices of applying concepts in judgment" (Williamson 2007, 220), she can content herself with condemning only those judgments which are plausibly sufficiently psychologically similar to judgments which experiments suggest to be problematic.[5] How similar is 'sufficiently' similar? That's a difficult question to answer, but it isn't one exclusive to the experimental critique – we face such issues when we, for instance, attempt to determine the consequences of psychological studies on memory for our treatment of eyewitness testimony.

The trickier issue is not the risk of skepticism *per se*, but rather the 'overlap' between philosophical and non-philosophical cognition – and it is here that the experimentalists' (and Williamson's) reliance on the standard categories of epistemology becomes problematic. Abandoning 'intuition' for a variety of more narrowly drawn mental categories might suffice to avoid global skepticism, but it does not fully avoid the overlap issue. Consider the fact that, although the cognitive processes underlying moral and logical cognition are plausibly different, those that underlie everyday knowledge attributions ("my mother knows that today is my birthday", e.g.) and attributions of knowledge made in the context of a philosophical thought experiment are plausibly largely the same.[6] Even a 'mere' rejection of epistemic intuitions, then,

[4] In fact, I think the needed categories will be much narrower than this – even moral judgment plausibly involves multiple psychologically and epistemologically distinct mental processes.

[5] This does tone down the ambition of experimental philosophy's negative program quite significantly. But the failure of the armchair-in-flames approach to methodological criticism doesn't at all show that philosophers should rest on their laurels – even a piecemeal criticism of problematic mental states has the potential to be quite devastating.

[6] I leave aside here questions of expertise; such issues would favour the experimentalist on this particular point, regardless.

would still invite the rather implausible view that ordinary, everyday knowledge attributions are unjustified. And *mutatis mutandis* for other cognitive processes upon which experimentalists might cast doubt. It is this issue – the *overlap* between the forms of cognition employed in everyday activity and in philosophical inquiry – that I take to be the primary challenge that Williamson's critiques present to experimentalist criticisms of philosophical method. And it is this aspect of the challenge that highlights, I'll argue, a need for a more fine-grained, pluralistic approach to epistemic states.

2.

One possible strategy for dealing with 'overlap' problem is to appeal to an epistemic difference between different practices which make use of intuition, rather than differences between the various sub-types of intuition itself. Jonathan Weinberg (2007) has argued for just such a strategy. Perception, Weinberg notes, is quite obviously fallible. Yet this fallibility is rendered unproblematic by our everyday epistemic practices:

> "our practices involving perception will by and large keep us from getting into too much epistemic trouble…[we] know a great deal about the circumstances and ways in which it is fallible and about what to do when we find ourselves in such circumstances" (Weinberg 2007, 324).

In Weinberg's terminology, our everyday perceptual practices are 'hopeful' – they allow us to detect and correct for the errors inherent in perception. By contrast, Weinberg suggests that the philosophical practices surrounding use of intuition are hope*less*. It is this hopelessness, rather

than the unreliability of intuition itself (which is plausibly not much greater than that of perception), which renders our philosophical practices worthy of criticism.

This is a move in the right direction, but it isn't enough to avoid the overlap problem. For it's not clear that philosophical practices surrounding use of intuition are actually more hopeless than the corresponding everyday practices surrounding use of intuition. It's not as though there are error mitigating procedures in ordinary cognition that we fail to avail ourselves of when philosophizing – if anything, our ordinary epistemic practices with intuition are plausibly less rigorous, and consequently *more* hopeless, than our philosophical practices with such. As far as I can tell, the only apparent feature of philosophical practice that is supposed to render it more problematic than everyday usage of intuition is the fact that philosophy sets "no constraints on how esoteric, unusual, far-fetched, or generally outlandish any given case may be" (Weinberg 2007, 321). Yet it's not at all clear that 'far-fetched' cases are (or should be) the sole targets of the experimentalist critique. Nor is it clear exactly how to draw the line between the 'esoteric' and the ordinary – many Gettier cases, for instance, are fairly mundane (see also Williamson (2016) on the difficulty of isolating the 'far-fetched'). If all this is more or less correct, then a condemnation of philosophical practices which employ intuition still threatens to generalize to much of ordinary cognition.

Recently, a somewhat different approach has been suggested both by Alexander and Weinberg (2014) and by myself (2011, 2015). This approach focuses on what I have called the 'epistemic demandingness' of philosophy as compared to ordinary cognition.[7] The core idea is that different tasks come with different epistemic requirements; an epistemic source that is 80%

[7] Alexander and Weinberg's term for the same phenomenon is 'error fragility'.

reliable, for instance, might be sufficiently reliable for some tasks but insufficiently reliable for others. In particular, if a single task requires multiple successful uses of the same source, the risk of failure quickly multiplies – an 80% reliable ability to distinguish crows from ravens results in only about an 11% chance of success if your task requires you to properly distinguish 10 such birds in a row.

The Gettier literature provides a perfect example of why philosophy may well be highly demanding, in the sense just described. Gettier judgments form a tiny minority of knowledge-attributions – an intuiter who erred *only* on Gettier cases would thus be a highly reliable knowledge attributor. Yet suppose this intuiter were to use the set of knowledge-attributions which she found intuitive as the sole basis for her theory of knowledge. The sliver of unreliability in her knowledge-attributions would, if undetected, quite likely lead our hypothetical intuiter to endorse the widely rejected JTB theory of knowledge. Conversely, it's quite possible that JTB is in fact the *correct* theory of knowledge, and that our generally-reliable ability to classify cases of knowledge has, due to a small stain of error, set us off on the wild-goose-chase that is Gettierology. In fact, given the heavy weight we philosophers give to single counterexample cases, it's plausible that just about any non-zero degree of unreliability brings with it the potential to massively disrupt successful theorizing.[8] Quite generally, in science as well as in philosophy, small amounts of error in a data source can lead to radically false theories – and if we want to minimize this risk, then we may need to implement especially stringent procedures to mitigate said errors.

[8] None of this, of course, demonstrates that any of our theories *are* so affected. But the likelihood of such a scenario may be significant enough to warrant concern about our methods - just as the likelihood of experimenter bias is significant enough to warrant obligatory double-blinding. Again, nothing about this requires any form of scepticism; if such an argument mirrors that of the skeptic (as might be suggested by discussion in Williamson (2007)), then so too does the argument for stringent methodological standards in science.

Suppose that this line of argument is on the right track. It suggests that we need to distinguish not only different practices which employ intuition, but also different *epistemic standards* which correspond to those practices. And this, I think, is the real heart of the issue – participants in the intuition debates have falsely assumed that a single epistemic standard suffices for both ordinary epistemological questions and for metaphilosophical questions about the quality of our methods. Yet it may well be that the epistemic standards operative in philosophical contexts are unusually high; perhaps issues like the epistemic demandingness of philosophical theorizing lead to a greater-than-usual need for vigilance against error. If this is indeed so, then a given source might be of sufficiently high quality for ordinary activity while being too error-prone for philosophical purposes; the overlap problem, at least potentially, dissolves.

There are, of course, accounts of knowledge which accept that epistemological standards may shift between contexts – contextualist and subject-sensitive invariantist accounts being the most familiar. But metaphilosophers aren't primarily concerned with the nature of knowledge; they are concerned with whether or not the practices employed by philosophers meet the epistemic standards appropriate to philosophical activity. And the familiar shifting-standards accounts simply don't directly address that issue. Suppose it to be true that philosophical activity involves higher epistemic standards. Even if a contextualist or related account *were* able to successfully accomodate such varying standards, questions like 'does intuition produce knowledge' would *still* be more or less irrelevant to questions about the viability of current philosophical methods; and plausibly, our ordinary epistemic vocabulary would *still* be inappropriate in metaphilosophical contexts.

To see why, suppose for the sake of argument that the only epistemic requirement that varies by context is the minimum reliability of the process by which a belief is formed. Suppose

further that subject-sensitive invariantism correctly captures this variation; in other words, suppose that the stakes are higher in philosophical contexts and that the standards for knowing are elevated in such high-stakes contexts. In particular, suppose that in ordinary contexts a reliability level of 80% suffices for knowledge; and suppose that, in philosophical contexts, a reliability level of 99% is required for the same. Now finally, suppose intuition to be 85% reliable. Is intuition sufficiently reliable to generate knowledge (in our imagined scenario)? Yes – it produces knowledge in at least some contexts. Is it sufficiently reliable to *generally* produce knowledge? Almost certainly, given the relative rarity of philosophical contexts compared to ordinary ones. Such questions are, however, unhelpful if our goal is a critical examination of our methods. The question we *should* be asking is as follows: is intuition sufficiently reliable to generate knowledge-in-a-philosophical-context? At which point, we may as well simply adopt a technical term – 'knowledge$_P$', say – to refer specifically to epistemic states that meet the standard we are interested in.

This is not to say that a subject-sensitive invariantist account *is* the right account of knowledge. In fact, I'll aim to argue that the epistemic norms underlying the various types of human cognitive activity are even more complex than the foregoing discussion has suggested; and that their complexity renders implausible any attempt to capture all the relevant phenomena under any kind of concept we would recognize as 'knowledge'. In that sense, the use of 'knowledge$_P$' and related terms is slightly misleading – the needed categories may not be straightforward analogues of standard epistemic categories, but rather wholly different types of epistemic state.

3.

We don't, I think, need to commit ourselves to any claim about epistemic demandingness in order to motivate the idea that epistemic standards vary dramatically between different forms of inquiry, and that epistemology's monolithic focus on knowledge is therefore problematic. We can simply consider our actual practices – for instance, the extraordinarily stringent methodological procedures that are considered obligatory in the sciences. Double-blinding is, of course, an obvious case here. The role of double-blinding in an experimental context is to eliminate potential biases resulting from the expectations of the experimenter or the subject – including such well-known phenomena as the placebo effect. The practice of double-blinding is not employed in every scientific activity, of course. But if the threat of bias is present, and if it is possible to blind the study without negative effects (such as ethical violations), then it is arguably obligatory that the study be conducted blind. It would be wholly appropriate to criticize a study that failed to adhere to this standard.

But note that double blinding is more or less *never* employed by ordinary cognizers during their everyday belief-forming activities. This is not because the relevant biases don't appear outside of scientific contexts. Ordinary cognition is rife with biases, including biases that result from the 'coloring' of observation by pre-existing expectations. We simply don't expect anyone to do much about it outside of the laboratory; at least, we certainly don't expect anyone to go to the lengths that scientists do to mitigate said biases. It is not as though ordinary cognizers, then, are typically failing to live up to their epistemic obligations. Instead, it is that our epistemic expectations – or at least, our methodological expectations – differ between these two types of epistemic activity.

Double blinding is only one example of many. Scientists are obligated to only rely on data that achieve statistical significance. They must control for confounding variables, where possible. They are obligated, in certain circumstances, to employ various measuring devices, videotaping protocols, multiple observers, and so forth to protect from the (minimal) errors that arise from ordinary visual observation. None of these procedures are ever required of ordinary cognizers – even when those cognizers are forming beliefs about the exact same propositions that are studied by science. Members of pre-scientific communities knew, for instance, all kinds of propositions about medicine, psychology, and so forth; they achieved this knowledge without use of the procedures that are now obligatory in those fields. To deny this would be to claim that, for instance, no one knew that any of our various medicines were effective before the placebo effect became recognized and standardly controlled for in the *mid twentieth century*.

The sciences, then, plausibly are subject to heightened epistemic standards which obligate practitioners to employ extra-stringent methods to control for bias and error – standards that exceed what is required for the attainment of knowledge in ordinary circumstances. I claim that the same is also true in philosophy – the epistemic standards to which professional philosophers ought to hold themselves are plausibly higher than the epistemic standards which must be met in order to know.

If epistemic standards really are elevated in fields like science and philosophy, then this prompts the following question: why should there be these different standards? If it is possible, for instance, to control for biases and reduce error, why shouldn't the methods for doing so be epistemically obligatory for everyone, in all contexts? Plausibly, it is simply because the practices that must be employed to live up to the stringent standards of professional inquiry are far too over-demanding for ordinary cognitive activity. Let's focus, for the moment, on the

sciences. Consider, for instance, how absurd it would be to criticize an ordinary cognizer for failing to run a double-blind controlled experiment before forming an ordinary causal generalization – say, the belief that the new office break policy has led to an increase in her co-workers' morale. The vast majority of non-scientists simply do not have the time, skills, or resources to subject their belief-forming to the rigorous methodological controls used in science. Employing scientific rigor in all instances of everyday belief formation would be inefficient, needlessly over-meticulous – obsessive, even. Nothing would ever get *done*.

It's common to argue that knowledge is implicated in various norms surrounding assertion, belief, and action. For instance, I should act on a given proposition only if I know that proposition to be true; I should believe a proposition only if I am in a position to know it. Suppose such arguments to be on the right track; knowledge, then, is a concept that is intimately linked with other aspects of our ordinary lives. If knowledge required ultra-stringent procedures for eliminating error, then, this would plausibly have extreme consequences for our practices of communication, practical reasoning, and so forth. If I cannot know p before subjecting p to rigorous experimental test, then this might imply that I cannot justifiedly act on p until such a test has been conducted. Again, that seems preposterously overdemanding; it seems absurd to think that ordinary individuals are epistemically required to invest so much time and effort before acting. Again, then, we have reason to think that ordinary folk are not epistemically required to meet the methodological standards which are obligatory in scientific activity.

Obviously, though, we do not consider it overdemanding to require that *scientists* subject their *scientific* activities to said methodological requirements. Now suppose we were to employ a straightforward contextualist or subject-sensitive invariantist account to try to capture these differing expectations. If one assumed a shifting-standards explanation of the differing

requirements – e.g., if one claimed that the standards for knowing are simply raised in scientific contexts – then this would suggest that the standards for believing, acting, and asserting would similarly be raised in said contexts. It's not clear to me that they are. Indeed, the methodological standards we've been discussing don't seem to place any limitations whatsoever on individual belief or action. A researcher who forms the belief that her theory is true does not thereby break the norms of science – so long as she dutifully continues subjecting her theory to the appropriate experimental tests. A researcher who takes her own experimental medication before fully verifying its effectiveness through experiment does not thereby break the norms of science - so long as she does not use her own case as data in her experiments. We might question whether said belief or action is warranted, but we would not accuse these hypothetical researchers of bad science: their actions do not make them poor researchers.[9] Contrast this with a case where a scientist fudges her experimental data, or fails to properly calibrate her instruments.

Scientific methodology is not concerned with regulating the beliefs and actions of individuals, except where those actions specifically violate required methodological procedures. The standards of science simply do not interface with the familiar knowledge norms we have been discussing. This makes perfect sense, for science is a community effort rather than an activity conducted by an isolated researcher. The community aspect of science, in fact, plausibly explains why we don't view the stringent methods of science as overdemanding in an experimental context. The limitations on time and resources that individuals face are not faced by the scientific community as a whole; if it takes several generations to confirm a hypothesis to the satisfaction of the scientific community, then so be it. The exemplars of knowledge upon

[9] The case of assertion is somewhat more complex. Certain assertions do seem to be prohibited by the methodological norms of science – at least in a publication context.

which traditional epistemology have been built are mental states of individuals; but a successful characterization of the epistemic features of group inquiry may well differ from traditional epistemology quite substantially.

I've argued that ordinary cognizers are not required to employ the strict bias-reducing and error-eliminating practices that are obligatory in scientific inquiry. But this should not be taken to imply that biases cannot undermine knowledge in ordinary circumstances; of course they can. If my belief that Smith is the most qualified candidate is based on a cherry-picked set of sources all of which back the view I favor, this might well prevent my belief from qualifying as knowledge (even if Smith is *in fact* the most qualified candidate). Indeed it *may* well be that a belief must be free of bias in order to count as knowledge (though I personally suspect that this is too strong a requirement). But notice that the requirements of science are higher still. It is not enough that a scientific study simply *be* free of bias; the experimenter must have taken explicit steps to *ensure* that the study is free of bias. A non-blinded study is still a flawed study, even if experimenter expectation did not *in fact* affect the experiment's results.

This might suggest to us that at least in the scientific case, what matters is not *knowing*, but knowing-that-one-knows. This might, in turn, enable an explanation of the apparent variation in methodological standards without requiring any sort of contextualist or subject-sensitive invariantist stance on knowledge. In fact Williamson has advocated just this sort of strategy for explaining our intuitions in standard cases of raised stakes (such as the well-known bank and airport cases):

> "[S]uppose that the agent in C knows p and the agent in C* knows p*, but the agent in C is in no position to know that she knows p and the agent in C* is in no position to know that she knows p*. Since stakes are higher in C* than in C, we as theorists may view the

failure of second-order knowledge in C* more sternly than its failure in C, and therefore regard p as appropriate in C but p* as inappropriate in C*" (Williamson 2005, 231).

In other words, when the stakes are higher an agent has more reason to 'double-check' on her knowledge – and we judge that such an agent has erred when they fail to do so. This need not entail that knowledge does not suffice for epistemically warranted action; locking one's door is sufficient for securing one's belongings from theft, but in some contexts we might nonetheless think it prudent to double check whether one *has* in fact met this standard. It would be open to Williamson to treat the variable standards among fields of inquiry in the same way – claiming that, in such fields, we think it is important to 'double check' that our data consists of genuinely known propositions.

Suppose, then, that we consider this as a possible framework for discussing metaphilosophical issues. Suppose, in other words, that we frame the experimentalist argument as follows: the flaws of intuition don't prevent it from being a source of knowledge, but they do prevent us from achieving second-order knowledge – that is, knowledge that we know. And, in philosophy, what we are really concerned with is knowing that we know. Note that, if this approach succeeds, there is no need for 'new' epistemic vocabulary, or for any departure from a classic knowledge-centric approach to epistemology. We could simply conduct metaphilosophical debate in terms of ordinary iterations of knowledge.

But unfortunately, iterations of knowledge aren't likely to provide a sufficiently fine-grained set of epistemic categories to capture the standards at issue. An examination of non-scientific fields of inquiry and their proprietary standards will, I think, demonstrate just how much variation we may need to countenance in order to capture the methodological requirements I've been discussing. This will be the focus of the next section.

4.

Science is not the only type of inquiry that plausibly operates with elevated epistemic and/or methodological standards. Consider some of the methodological obligations present in courtroom proceedings: 'leading' questions are forbidden, restrictions are put on admissible evidence, and verdicts must be held to be beyond reasonable doubt. Needless to say, these methodological constraints are not obligatory in most ordinary instances of belief-formation involving testimony or physical evidence. Or consider fact-checking expectations and standards of objectivity that hold in journalism. Or standards surrounding the use of written sources in history. And so on. Plausibly, nearly every academic and professional field has at least some methodological strictures aimed at improving epistemic position - strictures which they do not expect non-participants to either know about or hold to.

Indeed, some fields seem to have even higher epistemic standards – or, perhaps, higher epistemic aims - than the sciences. Consider the case of Goldbach's conjecture in mathematics, which states that every even integer greater than 2 is the sum of two primes. We have direct verification that this conjecture holds of all primes up to 4×10^{18}. Thus, we have tested the claim on *millions* of cases – many more confirming cases than most ordinary empirical experiments ever receive. Yet the conjecture lacks a proof; and mathematicians thus *continue* to work on the problem, and take it to be *wholly worthwhile to do so*. They will be satisfied with nothing less than either complete deductive certainty, or a demonstration that no proof is possible.

I find it plausible that we *know* Goldbach's conjecture to be true (or, more carefully, assuming that it is *in fact* true I find it plausible that our epistemic situation with respect to the conjecture suffices for knowledge). There are, of course, complicated issues surrounding the use of inductive evidence in mathematics. But similar cases suggest that we really should at least sometimes take ourselves to know in cases where proof is possible but absent. Consider the fact that several ancient cultures seem to have been aware of the truth of the Pythagorean theorem long before the notion of a proof was even first conceived – the well-known Plimpton 322 papyrus from c. 1800 bce Babylon, for instance, lists a large number of triples of integers satisfying $a^2+b^2=c^2$, suggesting that Babylonians possessed both a method for generating such triples and an understanding of their import. Presumably this understanding of the relationship between the sides of a right triangle was put to use by such cultures in practical contexts - say, in land surveying tasks. The fact that such actions seem epistemically warranted suggests that pre-Greek mathematics *is* genuine knowledge. Yet knowing is not enough for current mathematicians – Pythagoras's proof was not superfluous. Continued pursuit of proof in such cases (and others – consider the quest for a proof of Euclid's parallel postulate) indicates that mathematicians don't aim at mere knowledge, but at something much higher – perhaps complete certainty. Needless to say, nothing approaching this level of certainty is plausibly achievable in, say, science.

Even within a field, different standards presumably apply to different sub-fields, to different tasks, or even to different individual propositions.[10] The conditions under which a

[10] I take it this leaves open the possibility that the very same intuited proposition might be acceptable for use in one field but not another – especially if it is put to different tasks in those different fields. Williamson (2009) responds to Weinberg by noting that science makes use of intuitive epistemic judgments, when e.g. concluding that data renders a hypothesis probable. Yet it is possible that the use of intuition is acceptable for that task while being insufficiently rigorous for the purposes of, say, formulating an overall theory of knowledge.

scientist may assert "our experiment confirms our hypothesis" are obviously different from the conditions under which she may assert "our experimental group contained 235 subjects". Given this cornucopia of wildly different epistemic expectations, I think we should be a bit pessimistic about the idea that iterations of ordinary knowledge provide enough categories to capture this epistemic complexity. Do all fields aim for knowledge-of-knowledge? Are the standards in place in science exactly as rigorous as the standards in place in journalism, law, or history? What about mathematics – is there *any* iteration of knowledge that maps perfectly onto the deductive certainty that seems to play a central role in mathematics?

I can't fully consider such questions here. But the following seems, to me, to be a better overall characterization of the situation. There exist various methods and/or epistemic policies that have the potential to increase, to various degrees and plausibly along various dimensions, the quality of our epistemic position. When we are considering implementing some such rule, we are faced with a tradeoff: is the increase in epistemic position worth the non-epistemic costs (time, effort etc.) that come along with the adoption of the method? The answer to that question is likely to be different for an individual and for a group; it might be different for different fields, and for different tasks within a field. Stretching 'knowledge' and its iterations to capture all such acceptable trade-offs may be possible, but it's not clear why it should be necessary. As metaphilosophers we can, I think, simply adopt discipline-specific epistemic categories and not concern ourselves with the relation that such states may or may not have to knowledge.[11]

5.

[11] That's not to say that the question has no interest – but it is an epistemological question, not a methodological one.

Let's return, at long last, to the hypothesized elevated standards of philosophy. Does philosophy, like the other disciplines we've discussed, really aim higher than 'mere' knowledge? It is true that philosophy lacks the carefully articulated methodological requirements that operate within the sciences – there is no obvious analogue for practices such as double blinding, or even the use of measuring apparati. But it seems plausible that philosophy is nonetheless subject to epistemic standards that far exceed what we expect in ordinary cognition. In order for their work to be considered publishable, philosophers are expected to carefully define any crucial terms they employ in their arguments; they are expected to demonstrate familiarity with relevant literature; they are expected to anticipate and respond to obvious potential objections. They are subject to expectations of rigor, clarity, objectivity and the like that are more or less beyond the abilities of the majority of untrained, ordinary cognizers. Yet surely ordinary cognizers have knowledge of at least some propositions that have been, or could be, subject to philosophically rigorous debate. It is wholly implausible that no non-philosopher has ever known, for instance, any of the various ethical propositions, religious propositions, or so on over which philosophers have argued; not to mention potentially debatable propositions regarding the mind, beauty, causation, existence, and all the rest.

And even if we leave aside the question of whether philosophy *currently* employs elevated methodological standards, we can clearly make the case that it *should*. For note that the exact same considerations surrounding the elevated methodological standards in the sciences apply in philosophy, as well. Philosophy is a group effort, conducted by persons who are paid to spend large portions of their waking hours on philosophical inquiry. Constraints on available time, resources, skills, and so forth don't clearly apply to the philosophical effort as a whole. Supposing that at least some of the many mental processes that fall under the banner of

'intuitive' cognition do display biases of various sorts, it seems obvious that we ought to seriously concern ourselves with devising methods to mitigate such epistemic flaws. This is, essentially, the same argument that one might give in support of the use of double-blinding in the sciences.

It's not a very large step from there to the following claim: it is worthwhile for the philosophical community to divide its cognitive labor such that at least some members devote their time to empirical investigation of potential biases in philosophical cognition. It is also worthwhile for at least some members to devote their time towards consideration of possible mitigating methodological changes, or to impressing upon their fellows the potential pitfalls of overconfidence in a source whose epistemic credentials are not yet wholly established. It may well be that all current work in experimental philosophy is flawed; or that the findings thus far give us no good reason to accept the claim that 'intuitive' judgment displays bias. But our general awareness of the foibles inherent in nearly all areas of human cognition is, alone, enough reason to undertake the task of experimentally investigating the cognitive processes that we rely on as philosophers. One might even argue that doing so is obligatory – not for any individual philosopher, but for the philosophical community as a whole.

Here, then, is how I would frame an experimentalist critique of current philosophical methods. Intuition (or, better, most types of cognitive process that we lump together as 'intuitive') presumably generates knowledge – where this is understood as, say, the epistemic state that an individual must meet in order to assert or act in ordinary everyday circumstances. But empirical findings suggest that intuition has substantial flaws, and these flaws are incompatible with knowledge$_P$ – where this is understood as something like the epistemic standing we ought to aim for as professional philosophers (again, an even more careful argument

might distinguish even more epistemic standings, corresponding to the aims/obligations of different philosophical subfields, or required by different philosophical tasks, and so forth). This is, quite clearly, a more moderate position than 'intuition skepticism'. But it is also, I think, both a more subtle and a more plausible one, too.

Finally, I think it is an *obvious* position. It is, to my eyes, rather baffling that so much philosophical ink should have been spilled debating whether or experimental findings license a complete rejection of intuition. Obviously they do not. But just as obviously, they license a critique of intuition – one that cannot be easily formulated, however, if one insists on focusing solely on the presence or absence of knowledge, justification, or evidence. It is for that reason, then, that I have urged the introduction of novel epistemic concepts. Categories like the one picked out by 'knowledge$_P$' do not correspond to our pre-theoretic, 'folk' epistemological notions in any obvious way; but this, I think, should be no objection to their use. 'Knowledge$_P$' and its ilk are novel technical terms – but ones that delineate categories of genuine use to epistemologists and methodologists.[12] By means of such categories, we step out of the monistic, knowledge-centered epistemic framework which has obscured the mundanity of the methodological 'crisis' surrounding intuition – a crisis which exactly parallels the methodological challenges faced by scientists aiming to evade the biases of observation, or of courts aiming to evade the biases of memory. As with those challenges, skepticism is wholly beside the point.

[12] This need not imply that categories like Knowledge$_P$ denote 'natural kinds' or anything of the like; I claim merely that they reflect useful carvings of the multi-dimensional spectrum of possible epistemic standings.

Acknowledgement: The work described in this paper was fully supported by a Grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. LU 359613).

References

Adleberg, T., Thompson, M., & Nahmias, E. (2015). Do men and women have different philosophical intuitions? Further data. *Philosophical Psychology*, *28*(5), 615–641.

Alexander, J., & Weinberg, J. M. (2014). The "unreliability" of epistemic intuitions. In E. Machery & E. O'Neill (Eds.), *Current Controversies in Experimental Philosophy*. Routledge.

Alexander, J., Mallon, R., & Weinberg, J. (2010). Accentuate the negative. *Review of Philosophy and Psychology*, *1*(2), 297–314.

Bealer, G. (1992). The incoherence of empiricism. *Proceedings of the Aristotelian Society, Supplementary Volume*, 66, 99-138.

Bealer, G. (1996). On the possibility of philosophical knowledge. *Noûs,* 1-34.

Boyd, K., & Nagel, J. (2014). The reliability of epistemic intuitions. In E. Machery & E. O'Neill (Eds.), *Current controversies in experimental philosophy*. Routledge.

Buckwalter, W., & Stich, S. (2014). Gender and Philosophical Intuition. In J. Knobe & S. Nichols (Eds.), *Experimental Philosophy* (Vol. 2, pp. 307–346). Oxford: Oxford University Press.

Chudnoff, E. (2014). *Intuition*. Oxford: Oxford University Press.

Cullen, S. (2010). Survey-driven romanticism. *Review of Philosophy and Psychology*, 1(2), 275-296.

Deutsch, M. (2015). *The Myth of the Intuitive*. Cambridge, MA: MIT Press.

Feltz, A., & Cokely, E. (2008). The fragmented folk: more evidence of stable individual differences in moral judgments and folk intuitions. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1771–1776). Austin, TX: Cognitive Science Society.

Feltz, A., & Cokely, E. (2009). Do judgments about freedom and responsibility depend on who you are? Personality differences in intuitions about compatibilism and incompatibilism. *Consciousness and Cognition*, *18*(1), 342–350.

Goldman, A. (2007). Philosophical Intuitions: Their target, their source, and their epistemic status. *Grazer Philosophische Studien*, *74*, 1–26.

Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, *10*(2), 95–118.

Knobe, J., & Nichols, S. (2008). An experimental philosophy manifesto. In J. Knobe & S. Nichols (Eds.), *Experimental Philosophy* (pp. 3–14). New York: Oxford University Press.

Machery, E., Mallon, R., Nichols, S., & Stich, S. (2004). Semantics, Cross-Cultural Style. *Cognition*, *92*, B1–B12.

Machery, E., & O'Neill, E. (2014). Introduction: Experimental Philosophy: What is it good for? In E. Machery & E. O'Neill (Eds.), *Current Controversies in Experimental Philosophy*. Routledge.

Nado, J. (2011). *Intuition and Inquiry*. PhD thesis, Department of Philosophy, Rutgers University.

Nado, J. (2014). "Why Intuition?" *Philosophy and Phenomenological Research* 89: 15-41.

Nado, J. (2015). "Intuition, Philosophical Theorising, and the Threat of Scepticism." In E. Fischer and J. Collins, eds., *Experimental Philosophy, Rationalism, and Naturalism: Rethinking Philosophical Method*. Routledge. Pp. 204-221.

Nagel, J., San Juan, V., & Mar, R. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, *129*(3), 652–661.

Pust, J. (2000). *Intuitions as evidence*. Taylor & Francis.

Schnall, S., Haidt, J., Clore, G.L., & Jordan, A.H. (2008). "Disgust as embodied moral judgment," *Personality and Social Psychology Bulletin, 34*, 1096–1109.

Seyedsayamdost, H. (2015). On normativity and epistemic intuitions: Failure of replication. *Episteme*, *12*(01), 95–116.

Sosa, D. (2006). Scepticism about intuition. *Philosophy*, 81(04), 633-648.

Stich, S., & Tobia, K. (2016). Experimental Philosophy and Philosophical Tradition. In W. Buckwalter & J. Sytsma (Eds.), *Blackwell Companion to Experimental Philosophy*. Blackwell.

Swain, S., Alexander, J., & Weinberg, J. (2008). The instability of philosophical intuitions: running hot and cold on Truetemp. *Philosophy and Phenomenological Research, 76*(1), 138–155.

Tobia, K., Buckwalter, W., & Stich, S. (2013). Moral Intuitions: Are Philosophers Experts? *Philosophical Psychology*, *26*(5), 629–638.

Weinberg, J. (2007). How to challenge intuitions empirically without risking skepticism. *Midwest Studies in Philosophy, 31*, 318–343.

Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and Epistemic Intuitions. *Philosophical Topics*, *29*, 429–460.

Wheatley, T., & Haidt, J. (2005). Hypnotic Disgust Makes Moral Judgments More Severe. *Psychological Science*, *16*(10), 780–784.

Williamson, T. (2004). Philosophical "intuitions" and skepticism about judgment. *Dialectica*, *58*(1), 109–153.

Williamson, T. (2005). Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, *55*(219), 213-235.

Williamson, T. (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.

Williamson, T. (2009). Replies to Ichikawa, Martin and Weinberg. *Philosophical Studies*, *145*(3), 465-476.

Williamson, T. (2016). Philosophical Criticisms of Experimental Philosophy. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy*. Oxford: Wiley Blackwell.

Woolfolk, R. L. (2013), Experimental Philosophy: A Methodological Critique. *Metaphilosophy*, 44: 79–87.