



Knowledge Is Not Enough

Jennifer Nado

To cite this article: Jennifer Nado (2017): Knowledge Is Not Enough, Australasian Journal of Philosophy, DOI: [10.1080/00048402.2016.1273959](https://doi.org/10.1080/00048402.2016.1273959)

To link to this article: <http://dx.doi.org/10.1080/00048402.2016.1273959>



Published online: 15 Jan 2017.



Submit your article to this journal [↗](#)



Article views: 63



View related articles [↗](#)



View Crossmark data [↗](#)



Knowledge Is Not Enough

Jennifer Nado

Lingnan University

ABSTRACT

Discussions of the role of intuition in philosophical methodology typically proceed within the knowledge-centred framework of mainstream analytic epistemology. Either implicitly or explicitly, the primary questions in metaphilosophy frequently seem to revolve around *whether or not intuition is a source of justification, evidence, or knowledge*. I argue that this Standard Framework is inappropriate for methodological purposes: the epistemic standards that govern inquiry in philosophy are more stringent than the standards that govern everyday cognition. The experimentalist should instead view her criticisms as analogous to calls for the use of double-blinding in science.

ARTICLE HISTORY Received 8 October 2015; Revised 9 December 2016

KEYWORDS experimental philosophy; intuition; metaphilosophy

1. Introduction

For debates over philosophical methodology, the question of whether we know some philosophical proposition *p* is rarely of interest. The epistemic standards relevant to philosophical inquiry, I'll argue, are proprietary—and they are *more stringent* than the standards that must be met in order to know. This is not a contextualist claim about the semantics of 'know'; instead, it is a claim about the epistemic norms that hold in philosophical contexts.¹ Those norms differ substantially from those that govern ordinary belief formation, and thus a proposition can be known while failing to be admissible in philosophical debate.

Philosophy is not alone in this. Scientific inquiry demands similarly elevated standards, sometimes providing evidence for what is already known. These elevated standards seem to be at least implicitly recognized by the scientist—less so by the philosopher. Philosophers' recent discussions of philosophical methodology have largely proceeded within the traditional framework provided by mainstream analytic epistemology: either implicitly or explicitly, the primary questions in metaphilosophy seem to revolve around *whether or not intuition is a source of justification, evidence, or knowledge*. Others have recently questioned the metaphilosophical literature's focus on

¹ I will remain neutral on what an epistemic norm *is*, in a metaphysical sense (for instance, on whether norms can be naturalized).

intuition [Williamson 2007; Deutsch 2010; Cappelen 2012];² I will question its reliance on the knowledge-centred framework of standard epistemology.

2. The Standard Framework of Analytic Epistemology

Analytic epistemology is, rightly or wrongly, typically characterized as the study of knowledge. Most of its core concepts are linked directly or indirectly with knowledge: justification is thought to be (part of) what makes a belief a case of knowledge; evidence is characterized as that which provides justification [Feldman and Conee 1985], or identified with what is known [Williamson 2000]; (epistemic) rationality is said to involve conforming one's beliefs to one's evidence. This cluster of interconnected concepts provides what I'll call the 'Standard Framework' of traditional analytic epistemology.

Why are the states described by the Standard Framework of interest to philosophers? Arguably, it is in large part because of their normative implications. By uncovering the nature of justification, knowledge, and the like, philosophers articulate the circumstances under which an agent may reasonably form a belief, make an assertion, or employ a proposition in practical reasoning.³ The norms thus articulated are often assumed to be universal—everyone ought to aim their cognitive activities at the acquisition of knowledge, ought to form beliefs in accordance with their evidence, and so on.⁴ In practise, however, epistemological debate standardly focuses on individual knowers in ordinary 'everyday' contexts. This is not to say that philosophers discuss only beliefs formed in *ordinary circumstances*—epistemology is of course rife with fake barns and the like. But the epistemic principles we aim at uncovering via consideration of the not-so-ordinary cases are those that govern everyday belief.

By and large, this Standard Framework has been imported into debates in metaphilosophy. The epistemic claims of participants on both sides of the recent intuition literature are commonly broad and wholly general. Often, they explicitly employ Standard Framework concepts:

'Experimental evidence seems to point to the unsuitability of intuitions to serve as evidence at all' [Alexander and Weinberg 2007: 63].

'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].

Such examples could be multiplied—and, even when not explicit, the Standard Framework seems to lurk in the background.

'Evidence' may be the most commonly used of these standard concepts in metaphilosophical debate. Unfortunately, participants in the intuition debate almost universally

² I am sympathetic to that view—nonetheless, I have stated the arguments in this paper in terms of 'intuition' rather than do battle on two fronts simultaneously.

³ See, e.g., Williamson [2000], Hawthorne [2004], Stanley [2005], Sutton [2005], and Hawthorne and Stanley [2008] for various norms in which knowledge purportedly plays a central role. Alston [2005] provides numerous quotations characterizing the permissibility of belief formation in terms of justification.

⁴ Not all of epistemology conforms to this characterization—much work in formal epistemology, in particular, doesn't comfortably fit the picture that I've sketched. Participants in metaphilosophical debate, however, largely draw their concepts and background views from non-formal epistemology.

fail to articulate what, specifically, they mean by ‘evidence’,⁵ and there is certainly never any claim that the conception of evidence being employed is *specific* to philosophy. That is, nothing deters the reader from assuming that the relevant conception of evidence is one that applies to everyday belief formation.

‘Reliability’ is also frequently invoked. Tobia, Buckwalter, and Stich, for instance, claim that the Actor-Observer bias gives us ‘good reason to think that intuitions are unreliable’ [2013: 631]. Boyd and Nagel portray experimental philosophers as arguing that ‘we have empirical evidence for the unreliability of epistemic intuition’ [2014: 114].⁶ But, again, benchmarks for reliability are nowhere mentioned, leading us to suppose that the relevant standard is ‘reliable enough to justify belief’ or ‘reliable enough to ground knowledge’. The apparent assumption is that what is at issue is intuition’s status as a source of the standard epistemic goods of mainstream analytic epistemology.

Finally, defenders of intuition frequently portray themselves as combatting ‘intuition scepticism’. Elijah Chudnoff, for instance, argues against several forms of ‘scepticism’ about intuition, understood as ‘the view that intuition experiences do not justify us in believing propositions’ [2014: 98]. Joel Pust [2000], similarly, targets ‘explanationist scepticism’ about intuition. Although such anti-sceptical defences of intuition don’t always explicitly target experimental philosophers, experimentalist critiques are clearly within their scope. Consider, for instance, Timothy Williamson’s [2007] argument against intuition scepticism. According to him, the cognitive abilities underlying uses of ‘intuition’ in philosophy also underlie much of everyday cognition. Thus, any criticism of philosophical intuition will overgeneralize to, for instance, ordinary cases of concept application. In a later paper, Williamson unequivocally suggests that experimental philosophers’ ‘critique of reliance on philosophical intuitions will become a global scepticism’ [2016: 24].

But perhaps the concepts belonging to the Standard Framework of epistemology aren’t the right concepts for evaluating our methodological practices. It’s possible that philosophers aim not to attain mere ordinary knowledge, but rather to meet some higher epistemic standard. Knowledge, after all, is not the highest possible epistemic state—certainty, for example, is plausibly both more difficult to attain and more valuable than ‘mere’ knowledge. If philosophy does aim at a higher standard, then an experimental philosopher could recast her arguments in terms of alternate epistemological concepts—perhaps P-knowledge, P-justification, etc.—that correspond to this higher-than-knowledge aim.

Our hypothetical experimentalist could then quite readily admit that intuition is reliable, that it produces ordinary knowledge, and that its epistemic status is sufficiently good to permit its usage in everyday belief-forming contexts. She could do this while simultaneously denying that intuition generates P-knowledge, or that its epistemic status is high enough to render permissible our current philosophical practices. In short, she could claim that intuition is good, but not *good enough* for philosophy. Arguments

⁵ Williamson [2007] is a notable exception.

⁶ See Alexander and Weinberg [2014], though, for discussion of the difficulties in framing intuition debates in terms of reliability.

like Williamson's would become irrelevant; the threat of scepticism would immediately vanish.⁷ More generally, any argument defending the philosophical *status quo* via appeal to the role of intuition in ordinary cognition, to intuition's overall reliability, or to the plausible ability of intuitions to grant knowledge would simply fail to bear on the experimentalist challenge.

3. The Elevated Standards of Science and Philosophy

I'll argue that philosophy *does* aim at an elevated epistemic standard, and that the above argumentative move is therefore plausible. However, the case for elevated standards is easier to motivate within the sciences—particularly in branches of science that overlap with subjects of ordinary knowledge, such as psychology and medicine—and so this is where I will begin.⁸ I'll then argue that many of the considerations that make the elevated standards thesis plausible for science also make it plausible for philosophy. Finally, I'll discuss possible alternate explanations of these apparent elevated standards, such as contextualism or subject-sensitive invariantism. I claim that no currently available account precisely fits the phenomena that I'll be discussing; and that, even if some such account could be made to work, knowledge would *still* not be a useful category for purposes of assessing scientific or philosophical methodology. If this is so, then a substantial reorientation of the metaphilosophical literature may be in order.⁹

Let's begin with a case of ordinary generalization from experience. Suppose that an experienced schoolteacher, while driving home from work, reflects on the students whom she has taught during her career. After reasonably careful consideration, she comes to believe that students who mentioned being read to at home tended to have larger vocabularies. On this basis, she extrapolates to a *general* belief that students who are read to at home tend to have larger vocabularies. Suppose that her hypothesis is true. What is the epistemic status of the schoolteacher's belief? Likely we would need more information to determine whether she has knowledge; but plausibly at least *some* possible scenarios meeting this description count as knowledge, in the ordinary everyday sense.

However, *no* possible scenario meeting this description fulfils the stringent epistemic requirements of science. Suppose that a psychologist aims to run a study to test this same hypothesis. She will be expected to develop a consistent, readily quantifiable, measure of vocabulary size. She will be expected to attempt to control for possible confounding variables such as IQ or socioeconomic status. She will be required to show that her findings reach statistical significance. In short, her inquiry will be subject to a large number of quite strict norms—and should she break them, her study would become inadmissible in academic discourse.

⁷ Weinberg [2007] responds to the scepticism problem by suggesting that we focus on the *practices* of philosophers, rather than on features of intuition. Weinberg [2009] also notes that an epistemic weakness might affect, for example, a scientific practice more than an everyday practice. However, Weinberg stops short of suggesting that different *standards* are appropriate for the two sorts of practice, applying one standard ('hopefulness') to both. See Ichikawa [2012a] for critiques of Weinberg's strategy.

⁸ Science and philosophy are not alone here—as Jonathan Weinberg reminds me, law, journalism, and mathematics also plausibly employ proprietary epistemic standards.

⁹ See Ichikawa [2012b] for another perspective on why metaphilosophical questions ought to be distinguished from epistemological questions.

My intuition is that it is possible for someone to know that children who are read to at home tend to have larger vocabularies, despite meeting none of the methodological requirements of scientific inquiry. Similarly for other generalizations from experience which concern potential areas of scientific investigation. My intuitions here may well be more liberal than those of other philosophers. For what it's worth, however, I think that the more liberal intuition is fairly prevalent among the philosophically untrained. Laypersons often complain about scientific studies confirming 'what we already know'; and there is a vast number of 'common sense' claims that were widely believed long before they were scientifically verified. Examples include the fact that rich diets tend to cause obesity, or that sleep deprivation causes cognitive impairment. Did we know those facts before they were tested scientifically? Surely we did. Science often rigorously tests what is already known through common sense and a smattering of everyday observations. Occasionally, a common-sense belief turns out false, and therefore not known—but this does not show that the beliefs that *do* pass scientific scrutiny were not knowledge prior to being put to the test.

Of course, one should have a healthy distrust of intuition on cases like this. But there are also theoretical reasons for attributing knowledge in these non-scientific contexts. The bar for knowledge simply *needs* to be low if the man on the street is to possess much knowledge at all. Most contemporary epistemologists follow G.E. Moore in assuming that we fare reasonably well in matters epistemic, at least most of the time. If a theory implies that I don't know that Paris is the capital of France, or that I don't know that water is wet, then this is a strong mark against it. Whatever threshold we set for knowledge, then, must be such that people very regularly meet it.

Yet humans are monstrously imperfect belief-formers. We are prone to innumerable biases, inconsistencies, and so on. We are overconfident in our beliefs. We rely on heuristic 'shortcuts' such as availability and representativeness. For each of us, huge swathes of our beliefs have been formed, and subsequently maintained, under less-than-ideal conditions. Fortunately, a great many of those imperfectly formed beliefs are nonetheless true. Indeed, some of the heuristics that we employ are arguably adaptive, trading modest sacrifices in reliability for large savings in processing time and required cognitive resources.

Many of the true beliefs that result from these imperfect belief-forming processes standardly, and I think correctly, get treated as knowledge. Imagine how devastatingly poor our reserves of knowledge would be, were we to count as unknown any belief that had been tainted by confirmation bias, by bandwagon effects, or by the various influences of overconfidence, framing, cognitive dissonance, and all of the rest. If we agree with epistemologists about the prevalence of knowledge—and we should—then countless imperfectly formed beliefs will *have* to get classified as knowledge on any reasonably plausible account.

No such motivation seems relevant to science, however. Let 'S-knowledge' denote the epistemic state that results from inquiry that conforms to all appropriate scientific

norms,¹⁰ and that has as its content a true proposition.¹¹ It seems unproblematic to admit that no one was in possession of S-knowledge for most of human history. We could, similarly, introduce counterpart terms for the other concepts of the Standard Framework.¹² In fact, we may already have a concept like S-evidence—‘scientific evidence’ is often used by non-philosophers to refer to a type of evidence that has been obtained via scientifically sanctioned procedures.¹³ Meanwhile, we might use ‘O-knowledge’ to refer to ‘ordinary’ knowledge—that is, *mere* knowledge—and, *mutatis mutandis*, for justification and the like.

The methodological standards to which scientists hold themselves have not applied to any kind of cognitive activity until very, very recently; we cannot criticize Palaeolithic hunter-gatherers for failing to conform their beliefs to their (non-existent) S-evidence. Of course, this is partly because many of the relevant procedures simply did not exist. But, even now that we do have such methods, we do not and should not hold ordinary belief-formation to the sorts of standards used in the sciences.¹⁴ In ordinary epistemic activity, I would claim, we need strive no further than O-knowledge.

Consider, for example, the ubiquitous p-value used to measure statistical significance. If a given experimental finding has not been shown to meet the threshold for significance, it generally cannot be appealed to as evidence within a scientific context (thus, it is not S-evidence). But it would be absurd to require that ordinary cognizers similarly screen their evidence. Imagine if ordinary individuals, during their everyday cognitive activity, were epistemically required to run statistical analyses on their experiences before inferring causal relationships. That would be ludicrously over-demanding—it would slow belief-formation to a crawl.

Similarly, double-blind methodology is standard procedure in experiments where it can be applied: its role is to reduce effects arising from the subjects’ and the experimenters’ pre-conceptions or biases. It would be impractical, to say the least, to employ anything like double-blinding when making observations in ordinary contexts. It is not as though the aforementioned biases cannot arise in ordinary contexts; it’s simply that we don’t expect people to employ the sorts of bias-mitigating procedures that scientists use before forming their beliefs.

This asymmetry makes perfect sense, for there are constraints faced by individual epistemic agents that do not apply to the scientific endeavour as a whole. Individuals are under limitations with regard to time, available cognitive resources, access to

¹⁰ I haven’t yet given a characterization of ‘appropriate scientific norms’. Two options suggest themselves—the norms to which scientists currently expect themselves to conform, or the norms to which scientists ‘ought’ to conform. I’ll eventually plump for the latter, but for now S-knowledge will remain somewhat underdescribed.

¹¹ I want to leave aside questions of whether science aims at truth or, for example, successful prediction. The ‘truth’ clause is meant to keep S-knowledge fairly parallel to ordinary knowledge for expository purposes; I don’t, in fact, have a strong commitment to the particular features of S-knowledge.

¹² As an anonymous reviewer points out, it may be that the needed epistemic categories involve an even *greater* departure from the Standard concepts than a term like ‘S-knowledge’ suggests. We may not need ‘counterpart’ terms so much as an entirely separate epistemic vocabulary.

¹³ Note that it *may* be that, if p is S-evidence, then it must also be evidence (*simpliciter*). And similarly for S-knowledge. The claim is that the two categories differ, not that they fail to overlap.

¹⁴ I have been speaking as if there were a single set of norms used by all scientists for all purposes; but that is a gross oversimplification. Different norms/standards exist in different branches of science. Further, not every ‘scientific’ proposition is held to a similar standard. The conditions under which a scientist may reasonably claim ‘our experiment confirms our hypothesis’ obviously differ from the conditions under which she may claim ‘our experimental group contained 235 subjects.’ The norms surrounding the latter may be the norms governing everyday cognition.

information-gathering procedures, etc. If a given belief-forming process takes two weeks, but an agent must act in the near future, then the agent would be ill-advised to adopt that belief-forming process. When characterizing the norms that govern belief-formation for individuals in ordinary circumstances, all of these constraints may be relevant. By contrast, such constraints are not relevant to the norms of scientific inquiry. Any particular researcher may lack funds, or may face certain deadlines. But such constraints fall by the wayside when considering scientific inquiry generally—as an endeavour not of individual epistemic agents but of the entire scientific community, past, present, and future. It's not obvious why the norms, standards, and aims that apply to individuals and those that apply to a community over extended periods of time should be the same.

One salient difference here, then, is that scientific norms seem to apply in virtue of the fact that science is a community effort. Standard analytic epistemology focuses on the epistemic states of individuals—O-knowledge is a state that an individual possesses in virtue of certain features of her belief. There have been occasional protests against this perspective: one strand of social epistemology, for instance, investigates the conditions under which groups might qualify as believers or knowers [Gilbert 1987; Goldman 2004; Pettit 2003; Bird 2010]. That project, however, still employs concepts from the Standard Framework. My suggestion is that the epistemic states of interest in science are *not* those of the Standard Framework. I'll remain neutral on the question of whether S-knowledge is possessed by the scientific community as a whole or instead by individual members of that community; the existence of a community of inquirers seems clearly to be part of the explanation for the elevated standards of science, but the specifics of the social dimension are beyond the scope of this paper.

My claim, then, is that the conditions that must be met to achieve O-knowledge are very different from those that must be met to achieve S-knowledge. Analytic epistemologists have almost exclusively focused on the former, while largely neglecting the possibility that O-knowledge is only one theoretically interesting epistemic state among many. This is not to deny the importance of O-knowledge—a full understanding of the requirements of O-knowledge would help us to determine how we ought to form beliefs in ordinary epistemic circumstances. But it would not, in and of itself, provide guidance for the scientist. The norms surrounding scientific methodology, I claim, cannot be derived from consideration of the requirements that one must meet to attain O-knowledge.

Let's move now to philosophy. At first glance, philosophy appears to lack the formal requirements on method that the sciences impose. There are no direct analogues to practices like blinding or replication. However, we might make a case for some less formally articulated methodological expectations. When defending her views in print, a philosopher must carefully define any important concepts or terms that she employs; she must be familiar with relevant literature regarding her claims; she must make a reasonable effort to consider and respond to possible counterexamples to her account, along with possible alternate theories. It's plausible that there are further, less easily articulable expectations regarding rigour, clarity, objectivity, and so on. These expectations are not as obvious or as well-codified as those of the scientist; but it's clear that the expectations to which we hold our colleagues are higher than those we employ when, say, grading undergraduate papers.

So, under what conditions is it reasonable for a non-philosopher to form beliefs about philosophical propositions? Suppose that an intelligent student, without any

exposure to philosophy, reasons that the existence of evil is incompatible with a loving and all-powerful god, and comes to believe, on that basis, that such a god does not exist. The student has not thought of, or otherwise been exposed to, any of the classic responses to the problem of evil. Suppose, for the sake of argument, that God does not exist. Has the student done anything epistemically impermissible? Suppose that the student stops attending church, or asserts God's non-existence to his peers. The student's epistemic situation could be improved, certainly. But has the student met his epistemic obligations with regard to the proposition that God does not exist? Does the student O-know that God does not exist?

Suppose that an intelligent child, without any exposure to philosophy, reflects on the possibility of animal suffering and comes to believe, on that basis, that eating meat is wrong. The child has not thought of, or otherwise been exposed to, any further relevant arguments. The child stops eating meat on the basis of her belief, and asserts her belief to her parents. Suppose, for the sake of argument, that eating meat is wrong. Has the child erred epistemically, or failed to fulfil any epistemic duties? Does the child O-know that eating meat is wrong?

I hold that the agents in these cases possess O-knowledge, but fail to possess P-knowledge. And this, I think, is how we *must* treat these sorts of cases if we are to maintain a Moorean attitude regarding the prevalence of knowledge. After all, simply consider the vast range of propositions that might reasonably be deemed philosophical. Philosophy's subject matter has a considerable overlap with claims of 'common sense'; ordinary folk have numerous beliefs that have, at one point or another, come under philosophical scrutiny. But ordinary folk are, let's admit, not particularly good philosophers—they are typically bursting with conceptual confusions, inconsistent beliefs, and blurred distinctions. Are we prepared to say that no one other than a handful of philosophers has ever known any substantive philosophical propositions about morality, knowledge, existence, beauty? Of course, many non-philosophers have false beliefs on such topics. But, any way you slice it, there will inevitably be countless true 'philosophical' beliefs in the minds of the philosophically naïve. It is exceedingly implausible to suppose that none of them qualifies as knowledge; yet in a philosophical context those very same propositions demand careful, rigorous scrutiny. We should, then, grant the folk substantive O-knowledge of philosophical topics; but little to no P-knowledge.

Many of the considerations discussed earlier regarding science apply again in the case of philosophy. Ordinary cognizers are under constraints with regard to cognitive resources, time, access to education, and so on. Philosophy, by contrast, is a community effort performed by a group of individuals who are paid to devote a very high proportion of their cognitive activity to consideration of matters philosophical. The sorts of epistemic expectations that would be over-demanding when applied to ordinary individuals are reasonable when applied to philosophers. I think, then, that we ought to take seriously the possibility that philosophy is governed by norms that are more stringent than those governing ordinary belief formation.¹⁵

¹⁵ There are interesting questions to ask about how comparable science and philosophy are, and how strong an analogy between them can be. However, for current purposes all that is needed is the observation that both are inquiry-centred professions involving a community of inquirers with greater-than-normal access to time, resources, and the like. We *pay* professional inquirers to spend more time and energy on inquiry than the layperson could be expected to muster—we can thus demand of them increased epistemic returns in the form of higher-quality epistemic states.

4. Alternate Explanations

The existence of different epistemic expectations doesn't immediately show that metaphilosophers must look beyond the Standard Framework.¹⁶ Three types of alternate explanation come to mind, each of which might imply the importance of Standard Framework concepts for metaphilosophy. The first is that, despite appearances, there is only a single standard operative in the discussed cases. The second is that there are indeed separate standards, but that these can be accounted for by some form of contextualism or related account. The third again admits that there are separate standards, but claims that the standards of the scientist/philosopher do not reflect genuine epistemic norms. Let's explore these possibilities in turn. I'll focus on scientific norms in the discussion to come, purely because they are better-articulated than philosophical norms are.

One version of the first option is to deny that any possible scenarios like the 'schoolteacher' case qualify as knowledge—even O-knowledge—on the ground that there is a universal epistemic obligation to (say) take active steps to eliminate cognitive biases from one's belief-forming processes.¹⁷ Thus, this route denies the distinction between O-knowledge and P-knowledge. In light of the considerations rehearsed above, this would arguably amount to admitting that knowledge simply isn't as common as we thought. It would also amount to claiming that ordinary cognizers are subject to some quite demanding norms that might require extreme sacrifices of time, effort, and other resources. This strikes me as implausible. It's common to use over-demandingness as an objection to moral views like utilitarianism—shouldn't similar considerations apply in the epistemic case? Finally, one can't have it both ways—if we accept that most people don't know most propositions studied by science or philosophy, we can't really turn around and claim that experimental philosophers face a scepticism problem.

Alternatively, we might claim that the single standard is such that the teacher only counts as knowing if her belief-forming processes *would* pass the tests to which scientific studies are submitted. In other words, if one took all of the experience 'data' upon which the schoolteacher based her generalization and submitted it to statistical analysis, the resulting p-value would reach significance. Further, it must not be the case that her data were altered by any biases to which she might be subject, such as a tendency to overestimate the vocabularies of her female students.

This is less demanding than the previous option, but it would still imply that there is much less knowledge than most people generally assume there to be. I would hazard to guess that the majority of generalizations that people make are based on data that fail to reach statistical significance.¹⁸ I would also guess that a substantial portion of our observations are coloured by some amount of bias. More importantly, however, science *still* holds itself to a higher standard on such a proposal. It is not enough that a scientist's observations just *happen* to be free of bias—a scientist must take explicit, well-defined steps to *ensure* that her observations are free of bias.

A third possible single-standard explanation is that awareness of the human tendency to bias serves as a defeater. Non-scientists' casually formed causal generalizations

¹⁶ The Standard Framework might be perfectly suitable for ordinary epistemological questions. The proposal is that we supplement the framework, not that we abandon it.

¹⁷ I am assuming that the description given in the previous section is inconsistent with the required bias-eliminating activities.

¹⁸ Jonathan Weinberg notes that this likely includes this very generalization itself!

and so forth count as justified only because they aren't aware of the cognitive flaws for which scientists control. Since scientists do possess that awareness, they must adhere to special methodological requirements in order to achieve justification. Again, this move has implausible consequences. Scientists aren't just aware of biases while in the laboratory—that awareness persists. The 'defeater' proposal would imply that scientists are obligated to conform to the relevant methodological requirements even outside the lab. But it's clear they are not. Scientists likely put somewhat more effort into avoiding cognitive errors than does the average Joe, but it is unlikely that their everyday cognitive activities even approximate the rigour of scientific inquiry.

Taking a different tack, we might consider the hypothesis that some propositions are simply harder to know than others are, and that science is in the business of studying such propositions—this might explain the *appearance* of separate standards. After all, even assuming a universal standard for knowledge, some pieces of knowledge take more investigation to acquire. It takes only one observation to come to know that my cat has whiskers—it takes considerably more effort to learn that all cats have whiskers.¹⁹ There is something to this idea, for science is typically in the business of establishing generalizations rather than particular truths—not to mention claims regarding such intangibilia as laws of nature. Everyday belief-formation, by contrast, deals largely in particulars—'the laundry is finished', or 'we're out of milk.' Even in a scientific context, no double-blinding or p-value calculations would be required before such propositions were deemed admissible.

The fact that different propositions display different levels of what we might call 'epistemic demandingness' may well be part of the story. But it can't be the whole story, since the examples that we examined earlier (fatty foods cause obesity, etc.) were cases where both scientists and non-scientists formed beliefs about the exact same causal relationships. There is substantial overlap between the propositions that enter into ordinary cognition and those that are subject, in certain contexts, to the methodological norms of science. That goes double for philosophy. Any explanation that appeals to sharp divisions between the subject matters of science, philosophy, and everyday belief-formation is therefore implausible. It *is* plausible, however, that science's concern with 'hard propositions' may partially explain *why* scientists have implemented heightened standards.

Let's turn next to explanations that accept the existence of multiple standards, while still claiming that Standard Framework concepts are the relevant epistemological categories for metaphilosophy. One possible move would be to invoke a contextualist account of knowledge. Contextualism is a semantic thesis about the term 'knowledge'—the contextualist claims that knowledge claims express different propositions in different contexts. As a semantic thesis, contextualism gives us little insight into how the truth-value of a knowledge-attribution is meant to link up with the epistemic *norms* that we've been discussing. On standard contextualist accounts of knowledge, what determines the relevant standard is the *conversational* context of the *attributor*; so one might claim that the impermissibility of (for instance) invoking statistically non-significant findings stems from the fact that conversations by scientists about their own activities make possibilities for error salient, thereby raising standards. But this would

¹⁹ Of course, I could come to know the latter statement via testimony—but consider here an 'original' case of knowledge.

suggest that there could be conversations about scientific activities (perhaps by non-scientists) in which standards are not elevated—and in which it would be inappropriate to criticize a researcher’s appeal to non-significant data. Attributor-centred accounts do not capture what drives our judgments here; the standards to which scientific activity is bound hold regardless of attributor context. *Mutatis mutandis* for philosophy.^{20,21}

More promising would be a multiple standards account upon which the determinants of the active standard are subject factors, rather than attributor factors. On views of this type, the proposition expressed by an utterance ‘S knows that p’ does not vary. Instead, whether or not a subject knows p is sensitive to features of the subject that go beyond her belief in p and evidential assets regarding p. Standardly, the hypothesized extra features involve practical interests or stakes—if the consequences for error are serious, the standards for knowing are higher [Hawthorne 2004; Stanley 2005]. For some cases, this might make sense: we want to be absolutely sure that a medicine is safe before we prescribe it to the public. But much of science isn’t like this at all: the stakes surrounding investigation of the mating habits of red-tailed weasels, say, simply aren’t that high. A fan of practical-interest views might protest that the stakes for an individual researcher might well be quite high—discovering truths about weasel mating might make or break a young researcher’s career. But notice that stakes are generally different for, say, an untenured researcher than they are for an emeritus professor. Yet it’s extremely plausible that the same norms of inquiry apply to both.²²

Finally, the most promising sort of multi-standard explanation simply recognizes multiple grades of knowledge. Ernest Sosa [2001, 2007], for instance, distinguishes between ‘animal knowledge’ and ‘reflective knowledge’; Alvin Goldman [1999: 23] between ‘weak’ and ‘strong’ knowledge. Some philosophers view ‘understanding’ as a distinct variety of knowledge whose value exceeds mere knowing (e.g. Grimm [2006]).²³ And Stephen Hetherington [2001] characterizes knowledge as a scale, recognizing that one’s knowledge of a proposition might be ‘better’ or ‘worse’. But again, none of these views exactly captures the phenomena at issue. Sosa’s reflective knowledge is one’s knowing that one knows (for Sosa, apt belief that one aptly believes)—but clearly one might possess this without meeting the methodological standards of scientific inquiry. Goldman’s categories similarly fail to map on to ‘ordinary’ and ‘scientific’ grades of knowledge. Understanding, meanwhile, may indeed be one goal of scientific and philosophical inquiry; but methodological constraints such as double-blinding do not obviously aim at its production. They reduce the risk of error—but even immunity from error doesn’t automatically produce understanding. Hetherington’s view isn’t incompatible with the phenomena; but it also doesn’t provide ‘carvings’ of the scale of

²⁰ Another point against the contextualist explanation: as noted in note 14, even within a single scientific context different standards appear to apply to different sorts of propositions. It is difficult to see how to accommodate this within a standard contextualist account. (Thanks to an anonymous reviewer for this point.)

²¹ Contextualist accounts need not invoke conversational salience, as I have in my example here—however, the fundamental difficulty with a contextualist explanation is not the conversational aspect, but instead the attributor-centred nature of contextualism. Scientific standards are what they are—they do not vary depending upon who is attributing knowledge.

²² A subject-sensitive account which focused on features other than practical interests could conceivably be made to fit the phenomena—say, one where the relevant feature was the subject’s professional role as scientist/philosopher. Nonetheless, for reasons given below, distinctions corresponding to S-knowledge and P-knowledge will still be needed for purposes of methodological debate.

²³ Others (e.g. Kvanvig [2003]) deny that understanding is a form of knowing. For reasons detailed below, I think neither account fits the phenomena discussed here.

knowing corresponding to S-knowledge or P-knowledge. As I'll now argue, there's a real need for such categories.

Ultimately, it doesn't matter how we characterize knowledge (in the broad sense). If we accept that philosophers and scientists aim at a higher standard than is required for the lowest 'grade' of knowing, there is simply not much sense, from a methodological perspective, in debating whether a given procedure or source of evidence generates knowledge (in the broad sense). Early proponents of procedures like double-blinding were not epistemologists; they simply noted that biases negatively affected the epistemic value of scientific observations, and argued that steps should be taken to eliminate the problem. It would have been absurd for such proponents to claim that unblinded experiments produce no knowledge—this would imply that for most of human history we possessed no knowledge about, say, the effectiveness of various medicines. And it would have been equally absurd to respond to calls for double blinding by defending the on-balance reliability of perception or its overall tendency to produce knowledge.

It is, however, reasonable (indeed, crucial) to debate whether scientific methods generate S-knowledge. Suppose that the scientific community regularly employs some method M, and that the epistemic credentials of M come under suspicion. It's of no interest, methodologically, to ask whether M generally produces knowledge—given the infrequency of scientific investigation compared to ordinary belief-formation, M might *generally* produce knowledge, but only by producing 'low-grade' knowledge while consistently *failing* to produce states that meet the higher standards of science. So, the relevant epistemological category for methodological purposes is not knowledge but instead states meeting those higher standards: in other words, S-knowledge. Similar points hold for reliability, justification, and so forth.

The moral is general. Suppose that adopting some method M would improve a subject S's epistemic position with regard to proposition p, but would further have some non-epistemic cost in terms of cognitive effort, time, etc. Should S adopt M? In order to answer that question, one doesn't typically need to know the conditions under which S would know p. If the increase in epistemic position is worth the cost, then one should adopt M. Whether or not that's true might differ for the ordinary person and for the scientist or philosopher (and for the scientific or philosophical community as a whole). The nature of knowledge, then, is really rather beside the point.

A final consideration is relevant here. Scientific norms govern how a scientist is to construct her experiments, under what conditions the scientific community views a theory as well-confirmed, and so forth. But they do not as obviously govern individual *belief-formation*. Suppose that a particular theory predicts the occurrence of a certain chemical reaction under such-and-so circumstances, but that no experiment has yet been performed to test this prediction. Suppose that a scientist is now in the initial stages of designing such an experiment. She is optimistic—on the basis of several previous successful predictions that the theory has produced, she forms the belief that the chemical reaction will occur as predicted. One might debate whether her belief is justified, but she does not *break the norms of scientific inquiry* by forming even a quite strongly held conviction that her experiment will succeed—her belief does not, for

instance, render her study inadmissible as S-evidence. The norms of scientific inquiry just don't cover this type of epistemic activity.²⁴ If this is right, then a 'multiple-grades-of-knowing' account may slightly oversimplify the situation—for, if S-knowledge were simply a 'higher' grade of knowing, we would expect the normative implications that arise from it to govern permissible belief formation in scientific contexts. It seems that they do not.

This brings us to a third possibility. If scientific norms aren't governing individual belief-formation, in what sense are we speaking of *epistemic* norms at all? If we take an epistemic norm to be a norm describing how one ought to believe, then the above considerations would indeed bar scientific norms from the category of epistemic norms. However, the above test would also deem the knowledge norm of assertion to be non-epistemic. And the norms under discussion don't plausibly seem moral or practical. But there is another possibility: perhaps the standards merely reflect conventions of the discipline. After all, the common .05% measurement for significance, for example, is largely arbitrary. However, there surely must be *something* more going on than brute convention. Scientific norms—unlike, say, norms of etiquette—plausibly have force even for communities who have not adopted them. Were we to encounter a group of extra-terrestrial scientists who set their p-value threshold at .08, we might deem that acceptable; but we would likely criticize alien researchers who, say, drew conclusions about human behaviour on the basis of a sample of test subjects abducted solely from American Midwest farming communities.

Ultimately, how we sort scientific norms makes little difference for current purposes. It's even possible that these phenomena merely indicate that the norm of assertion is different in scientific and philosophical contexts.²⁵ The upshot is still the same: the concepts of the Standard Framework are of little use in methodological debates, and there's simply no plausible sceptical problem to be dealt with. Scientists can and do appeal to potential biases involved in perception in order to criticize current scientific practices, without courting scepticism; what prevents experimental philosophers from doing the same for intuition?

5. Conclusion

I've argued that the standard concepts and questions of epistemology are too broad and too general to be useful when discussing how scientists ought to conduct inquiry: they fail to distinguish between everyday and scientific contexts. As weaknesses inherent in perception have been uncovered, scientists have not enmired themselves in debates over whether perception really leads to knowledge—they have simply implemented

²⁴ They also don't seem to govern most cases of individual action. The norms of science don't, for example, forbid a scientist from taking fish oil in an attempt to cure her cancer even though she has not yet obtained enough experimental evidence to verify her hypothesis that such a treatment has efficacy. My view does not, then, conflict with the views that knowledge epistemically suffices for belief and action.

²⁵ *Prima facie*, my view is inconsistent with the knowledge norm of assertion. One possibility, as just mentioned, is that there is a wholly separate norm of assertion in scientific/philosophical contexts. Another is that, on the correct characterization of the norm, knowledge is only necessary and not sufficient. This view has some support—see, for example, Lackey [2011]. Finally, given the discussion above, it may be that scientific norms limiting assertion are non-epistemic—and thus that they no more challenge the knowledge norm of assertion than do, say, politeness norms that forbid certain assertions about my colleague's weight.

appropriate correctives, including the use of specialized measuring instruments, multiple observers, eye-tracking technology, and so forth.

I've also, I hope, made at least a preliminary case that philosophy also employs proprietary norms, rendering obligatory certain levels of clarity, rigour of argumentation, and so forth—although a precise characterization of such norms is difficult to offer. Our current standards are thus more stringent than those that govern everyday cognition about matters philosophical. Yet, plausibly, the standards that we *should* be adopting are higher still. Experimental philosophy has given us reason to suspect that many types of intuition may be sensitive to cultural background, emotional state, framing effects, and so on. Yet philosophers have not developed any procedures to minimize the impact of such biases. We are not even taking the problem seriously enough to attempt to develop methodological corrections.

It seems to me that experimentalist criticisms of intuition are best framed as follows. There is an epistemological standard that we as philosophers *should* meet but currently do not—we should eliminate sources of bias and unreliability from our P-evidential resources. Intuition might be a reliable source of O-knowledge, but, so long as its biases are uncontrolled, it does not generate P-knowledge. If it is possible to control for such biases, even partially, then methods for doing so must become part of standard philosophical procedure. If it is not, we must search for alternate sources of P-evidence. Nothing in such an argument threatens us with scepticism. One cannot respond to such an argument by asserting intuition's status as a source of justification, evidence, or knowledge. The epistemology of intuition and the methodological problems surrounding intuition are best kept separate; metaphilosophers can and should look beyond the Standard Framework.²⁶

Funding

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

- Alexander, J. and J.M. Weinberg 2007. Analytic Epistemology and Experimental Philosophy, *Philosophy Compass* 2/1: 56–80.
- Alexander, J. and J.M. Weinberg 2014. The 'Unreliability' of Epistemic Intuitions, in *Current Controversies in Experimental Philosophy*, ed. E. Machery and E. O'Neill, New York: Routledge: 128–45.
- Alexander, J., R. Mallon, and J.M. Weinberg 2010. Accentuate the Negative, *Review of Philosophy and Psychology* 1/2: 297–314.
- Alston, W.P. 2005. *Beyond 'Justification': Dimensions of Epistemic Evaluation*, Ithaca, NY: Cornell University Press.
- Bealer, G. 1992. The Incoherence of Empiricism, *Proceedings of the Aristotelian Society, Supp. Vol.* 66/1: 99–138.
- Bealer, G. 1996. On the Possibility of Philosophical Knowledge, *Philosophical Perspectives* 10: 1–34.
- Bird, A. 2010. Social Knowing: The Social Sense of 'Scientific Knowledge', *Philosophical Perspectives* 24/1: 23–56.
- Boyd, K. and J. Nagel 2014. The Reliability of Epistemic Intuitions, in *Current Controversies in Experimental Philosophy*, ed. E. Machery and E. O'Neill, New York: Routledge: 109–27.
- Cappelen, H. 2012. *Philosophy without Intuitions*, Oxford: Oxford University Press.

²⁶ Thanks to Jonathan Jenkins Ichikawa and Jonathan Weinberg for comments on previous drafts.

- Chudnoff, E. 2014. *Intuition*, Oxford: Oxford University Press.
- Deutsch, M. 2010. Intuitions, Counter-Examples, and Experimental Philosophy, *Review of Philosophy and Psychology* 1/3: 447–60.
- Feldman, R. and E. Conee, E. 1985. Evidentialism, *Philosophical Studies* 48/1: 15–34.
- Gilbert, M. 1987. Modelling Collective Belief, *Synthese* 73/1: 185–204.
- Goldman, A.I. 1999. *Knowledge in a Social World*, Oxford: Clarendon Press.
- Goldman, A.I. 2004. Group Knowledge versus Group Rationality: Two Approaches to Social Epistemology, *Episteme* 1/1: 11–22.
- Grimm, S.R. 2006. Is Understanding a Species of Knowledge? *The British Journal for the Philosophy of Science* 57/3: 515–35.
- Hawthorne, J. 2004. *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Hawthorne, J. and J. Stanley 2008. Knowledge and Action, *The Journal of Philosophy* 105/10: 571–90.
- Hetherington, S. 2001. *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*, Oxford: Clarendon Press.
- Hetherington, S. 2011. *How to Know: A Practicalist Conception of Knowledge*, Malden, MA: Wiley-Blackwell.
- Ichikawa, J.J. 2012a. Experimentalist Pressure Against Traditional Methodology, *Philosophical Psychology* 25/5: 743–65.
- Ichikawa, J.J. 2012b. Experimental Philosophy and Apriority, in *The A Priori in Philosophy*, ed. Albert Casullo and Joshua C. Thurow, Oxford: Oxford University Press: 45–66.
- Kvanvig, J.L. 2003. *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lackey, J. 2011. Assertion and Isolated Second-Hand Knowledge, in *Assertion: New Philosophical Essays*, ed. Jessica Brown and Herman Cappelen, Oxford: Oxford University Press: 251–76.
- Pettit, P. 2003. Groups with Minds of Their Own, in *Socializing Metaphysics: The Nature of Social Reality*, ed. F.F. Schmitt, Lanham, MD: Rowman and Littlefield: 167–93.
- Pust, J. 2000. *Intuitions as Evidence*, New York: Garland.
- Sosa, E. 2001. Human Knowledge, Animal and Reflective, *Philosophical Studies* 106/3: 193–6.
- Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1*, Oxford: Oxford University Press.
- Stanley, J. 2005. *Knowledge and Practical Interests*, Oxford: Clarendon Press.
- Sutton, J. 2005. Stick to What You Know, *Noûs* 39/3: 359–96.
- Swain, S., J. Alexander, and J.M. Weinberg 2008. The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp, *Philosophy and Phenomenological Research* 76/1: 138–55.
- Tobia, K., W. Buckwalter, and S. Stich 2013. Moral Intuitions: Are Philosophers Experts? *Philosophical Psychology* 26/5: 629–38.
- Weinberg, J., S. Nichols, and S. Stich 2001. Normativity and Epistemic Intuitions, *Philosophical Topics* 29/1–2: 429–60.
- Weinberg, J.M. 2007. How to Challenge Intuitions Empirically Without Risking Scepticism, *Midwest Studies in Philosophy* 31/1: 318–43.
- Weinberg, J.M. 2009. On Doing Better, Experimental-Style, *Philosophical Studies* 145/3: 455–64.
- Williamson, T. 2000. *Knowledge and its Limits*, Oxford: Oxford University Press.
- Williamson, T. 2007. *The Philosophy of Philosophy*, Oxford: Blackwell.
- Williamson, T. 2016. Philosophical Criticisms of Experimental Philosophy, in *A Companion to Experimental Philosophy*, ed. J. Sytsma and W. Buckwalter, Malden, MA: Wiley-Blackwell: 22–36.