

Who wants to know?

Jennifer Nado

University of Hong Kong

The year is 1947. A normal, well trained doctor examines a patient who has come to his office complaining of a mild fever and headache. After ruling out any serious illness, the doctor recommends that the patient take a tablet of aspirin three times a day to reduce the fever and control the pain, and that she check back in if symptoms continue. Does the doctor know that aspirin is effective for the treatment of mild pains and fevers?

It's not a trick question. I'm expecting you to say 'yes'. Barring some severe deficit in their education, any doctor in 1947 would most certainly know the basic properties of aspirin. 1947, however, is one year before the first published randomized controlled trial of a medication (Bhatt 2010); it is also eight years before the publication of Henry Beecher's "The Powerful Placebo", which popularized the idea of placebo-controlled studies (Beecher 1955).¹ Modern experimental methods in clinical trials are in fact a surprisingly recent phenomenon; a 1951 analysis of 100 then-current trials found that 45% lacked even a basic control group, and another 18% employed inadequate controls (Ross 1951). Whatever evidence doctors had for the effectiveness of aspirin in 1947, then, would almost certainly not pass muster for (say) approval of a treatment for use in the United States by the current Food and Drug Administration.

Notice what this implies. When medical researchers investigate the properties of a new drug, the use of double-blinded, randomized, placebo-controlled trials are in many cases treated as obligatory² - quite plausibly, *epistemically* obligatory, though we'll return to that claim later. But our hypothetical doctor shows that it is clearly possible to know propositions of the form 'drug x is effective for treatment of symptom y' in the absence of such obligations having being met. I'd go further still - I'm quite comfortable with the idea that Hippocrates knew the medicinal properties of willow bark (the botanical source of aspirin) in the 5th century b.c.e., more than a millennium before the emergence of modern medicine. Again, whatever evidence Hippocrates possessed is woefully insufficient from the perspective of, say, the methodological standards demanded by modern medical journals.

If all this is correct, then it seems to me to suggest the following claim: current-day medical researchers don't aim their inquiries at the production of knowledge. Not as an ultimate goal, and not even as a means to an ultimate goal of true belief. Instead, medical researchers appear to aim for a more stringent, 'higher' epistemic state - one that reduces the chance of error far below the normal thresholds required for knowledge. More carefully, then, these researchers do not aim at the production of knowledge *qua* knowledge, but instead at a brand of ultra-high-quality knowledge - five-star knowledge, if you will. A bit of reflection will indicate that a similar claim could be made for many other branches of inquiry that involve proprietary, specialized methodological procedures, such as the

¹ The existence of placebos had by then been recognized for some time, but they were primarily used for non-research purposes - e.g., to mollify hypochondriacal patients.

² There are of course exceptions under circumstances that would prevent use of RCT's - for instance, if withholding treatment from a control group would be unethical.

various sciences, mathematics, law, journalism, and quite plausibly philosophy. For most academic disciplines and other inquiry-centered professional fields, knowledge *simpliciter* is simply not what inquirers want.

That's what I'll argue, at least. Obviously, there are other reactions one might have to the apparent ability of doctors to know medical truths prior to the mid-twentieth century. The first thought that's likely to come to mind would be to account for the apparently elevated standards of the modern medical community via one of the many available shifting-standards accounts of knowledge – some variety of contextualism or subject-sensitive invariantism, presumably. Or one could find some sort of argument to deny that the 1947 doctor and the modern medical community face different epistemic requirements, despite appearances. But I hope to convince you that such moves are much less plausible than they might initially appear. A much simpler option – and one which I see no real reason to resist – is to embrace a form of epistemic pluralism. Knowledge is not the be-all and end-all of human cognitive activity, but only one category of epistemic good among many. And for many epistemic undertakings, knowledge alone is just not good enough.

1. Our knowledge-centered epistemology

I think it would be fairly uncontroversial to claim that knowledge has been, and continues to be, the most central epistemic concept within traditional, non-formal epistemology. Indeed, 'epistemology' itself is most commonly defined as the study of knowledge. The focus is likely rooted in Plato, among others, but is most fully evident in contemporary epistemology; as any undergraduate philosophy major knows, 20th century epistemology was littered with endless attempts to analyze the concept of knowledge, particularly in the years immediately following Gettier's counterexamples to the classic JTB analysis.³ Of course, there was also much attention paid to such notions as justification, but justification itself has generally been considered to inherit its importance from its role in the analysis of knowledge. In these early years of the 21st century, there has been increasing skepticism about the prospects for a successful analysis of knowledge; yet we've also seen the growth of 'knowledge-first' epistemology. Knowledge continues to hold center stage.

1.1 The value of knowledge

There have been numerous attempts to elucidate just why knowledge should seem so important, and so central. Typically, these attempts focus on the question of why knowledge should be more valuable than mere true belief. The *locus classicus* for such accounts occurs in Plato's *Meno*, in which Plato first raises the problem by noting that true belief appears just as good as knowledge for achieving, say, the goal of finding one's way to the city of Larissa. Why, then, should we value knowledge as highly as we do? Plato's answer is that knowledge, unlike true belief, is 'tethered' – less apt to disappear when one is questioned or presented with reasons to doubt.

³ Here and throughout the paper, I restrict my attention to non-formal epistemology.

Contemporary responses to the ‘Meno problem’ – and to the more general question of why knowledge deserves special attention⁴ - are varied. Virtue epistemologists have suggested that their views can account for the value of knowledge by casting knowledge as a sort of cognitive achievement arising from epistemic virtue; this sort of achievement, in turn, is held to be intrinsically valuable (Zagzebski 2003, Greco 2003). Reliabilists like Alvin Goldman have located the value of knowledge in the fact that the reliability that accompanies knowledge makes our future beliefs likely to be true (Goldman and Olsson 2009). Edward Craig (1991) suggests that the concept of knowledge developed in order to help us to identify good sources of testimony. Timothy Williamson (2000) claims that knowledge holds interest due to its status as the most general factive mental state.

Another popular way to account for the centrality of the knowledge concept in epistemology is to claim that it plays a role in certain distinctively epistemic norms. Williamson (2000), for instance, provides extensive argument for the claim that knowledge is the norm of assertion, in the sense that one should only assert *p* if one knows that *p*. Fans of subject-sensitive invariantism have suggested that knowledge is the norm of action, in the sense that one should only act upon *p* if one knows that *p* (Hawthorne 2004, Stanley 2005, Hawthorne and Stanley 2008). Either of these suggestions, if true, would clearly show knowledge to merit philosophical attention.

Finally, a few authors have suggested that knowledge is also the “aim” of belief; we might think of this as claiming that a belief ‘gets it right’ if and only if it qualifies as knowledge (Peacocke 1999, Williamson 2000). More traditionally, the aim of belief has been seen as not knowledge, but mere truth (see e.g. Velleman 2000). Yet even if one accepts the latter claim, knowledge may still occupy a quite central role – Ralph Wedgwood, for instance, claims that as inquirers we aim at knowledge as a *means* for achieving the ultimate aim of true belief, noting that “we almost never aim to have true belief without at the same time aiming to know” (Wedgwood 2002, 289).

Nonetheless, a few contrary souls deny that knowledge has any special value that justifies its central role in epistemology. Jonathan Kvanvig (2003), for instance, argues that the value of knowledge does not exceed that of its parts. He argues that epistemological theorizing would do better to focus on the epistemic state of *understanding* – noting that this is not a mere species of knowledge, since one may understand without knowing and vice versa. Understanding, Kvanvig argues, does have unique value, and is thus more worthy of philosophical attention.

Though I won’t argue the case here, I’m at least a little inclined to think that knowledge is a reasonably useful concept for epistemology, whether due to its intrinsic value or for some other reason. Many of the arguments discussed above seem to me to be quite plausible – knowledge does seem to play some sort of role in regulating proper assertion and practical decision making, for instance, and the concept of knowledge also quite plausibly plays a crucial role in helping us to identify reliable sources of testimony. Or at least, all this seems to me to be true in normal everyday contexts, which are the contexts that epistemologists have standardly focused on (when they are not considering bizarre skeptical scenarios, that is). Whether or not the considerations above show that knowledge has some distinctive theoretical value, what’s important for my purposes is this: none of them rules out the possibility of *other* types of epistemic state which *also* have epistemic interest or value - possibly even

⁴ Pritchard (2007) helpfully distinguishes between the ‘primary’ value problem posed in the Meno and a ‘secondary’ value problem – why should knowledge have more value than any proper subset of its parts, such as justified true belief?

more interest or value than knowledge. Understanding may be one of these; but it's certainly not the only possible non-knowledge state of theoretical interest. I'll be aiming to convince you that there may well be many.

1.2 Epistemic states higher than knowing

No matter one's view on the value of knowing, it seems clear enough that knowledge is not the highest possible epistemic standing one might achieve (nor the lowest, for that matter). Though the traditional view of knowledge held it to require complete certainty or infallibility (for instance in Plato and Descartes), contemporary epistemologists typically concede that such requirements lead to skepticism. We possess vanishingly few beliefs, if any, that meet such stringent criteria. Most epistemologists seem to take the following as constraint on a successful account of knowledge: one's account must not be so stringent that it implies that none of us possess any knowledge at all.

So if a theory of knowledge leads to skepticism, this will be considered by many to be fatal to the theory. But would it be any better if a theory of knowledge implied that the vast majority of humans *almost* never achieve knowledge? Most of us would not be satisfied with such an account. We want our theory of knowledge to do justice to the Moorean sentiment that most of us know quite a lot – we know our own names, we know where we left our keys (usually), we know that water is wet, we know that the sky is blue. So the real constraint seems to be this: a theory of knowledge must not be so stringent as to rule out a significant proportion of those everyday true beliefs which we commonly class as known. Knowledge, then, will presumably need to be compatible with epistemic standings that fall quite a bit short of certainty.

There's a further reason to think that the bar for knowing must be fairly low. If knowledge does play a role in norms governing assertion and action, then a high-standards view of knowledge will imply that most of our assertions and actions are unwarranted (Hawthorne 2004). This is implausible, to put it mildly. We're inclined to think that large proportions of our practical reasoning and of our acts of assertion are perfectly acceptable. Note too that given the time-sensitive nature of much of our practical reasoning, it must be possible in many cases to get at knowledge rather quickly and easily. We are not generally at leisure to gather all possible evidence that might be relevant to determining the facts – we quite frequently must act on less than total information.

All this suggests that knowledge is a less demanding state than total certainty. But those two types of state only scratch the surface of the possible carvings of epistemic space. Understanding, of course, has already been mentioned. But why stop there? Suppose, as many epistemologists do, that knowing does not require knowing that one knows. Then knowing that one knows that P is presumably a more demanding (and presumably a more valuable) state than merely knowing that P; yet it's arguable that one can know that one knows P without total certainty that P. Or, consider the plausible claim that *some* level of reliability in the belief-generating process is necessary for knowing (sufficiency being a separate question). So long as the required reliability is less than 100%, we can speak of categories of epistemic state which involve a higher level of reliability than is strictly required for knowing, while falling short of utter certainty.⁵ The possibilities are nigh-endless - depending on one's views on closure,

⁵ Note that if you object to reliability being even a necessary condition for knowing, then this is even more obviously true.

safety, sensitivity, and all the rest, one can define up all sorts of epistemic state types that exceed the requirements for knowing by demanding whatever features knowledge fails to necessitate.⁶

So let's put aside the question of whether and why knowledge is more valuable than mere true belief; I'm more interested in asking the question of why knowledge should be so *uniquely* of interest given the all the potential epistemic standings that exceed merely truly believing. My claim will be that knowledge is simply not the only such state of interest. Knowledge may be the concept which regulates appropriate epistemic behavior in most 'ordinary' circumstances, but it does not reflect the elevated epistemic standards that we have implemented in specialized, professional fields of inquiry. Within these fields, inquirers appear to aim at epistemic states that exceed knowing. In the next section, I'll begin to explore this idea via the framework of epistemic normativity.

2. Obligations in Ethics and in Epistemology

Discussions of the value of knowledge naturally tend to invoke a notion of epistemic normativity – a notion of 'oughtness' which is parallel to, but at least *prima facie* distinct from, the notion of normativity familiar from moral theory. A notion of epistemic normativity is of course directly implied by those who claim knowledge to be the norm of assertion or action, or the aim of belief. But quite generally, philosophers and laypersons alike are comfortable making normative evaluations regarding matters epistemological. We will say that an agent *ought* to believe such and so, or to infer in such and so ways. We will say that someone made a *mistake* in reasoning. We quite regularly evaluate others' epistemic performance, and distribute praise and blame on the basis of said evaluation.

2.1 Characterizing epistemic normativity

I will be making use of various normative terms in what follows, in particular the notions of epistemic obligation, epistemic permission, and the oft-neglected epistemic supererogation. I'll be claiming, in short, that professional inquirers are under heightened epistemic obligations. Before I do so, however, a number of clarifications and caveats regarding epistemic normativity are in order. First: I don't intend to commit myself, here, to any particular position on the metaphysical status of norms. As a card-carrying naturalist, I'd prefer to avoid any characterization of norms that would leave them irreducible to more familiar ontic categories – fortunately, there are many naturalistically acceptable interpretations of normative talk available. One possibility would be to cast the norms I'll be discussing as merely expressing claims about the epistemic expectations individuals in our society have of one another. Another would be to express them as mere facts about means-ends relationships – as in, 'if an agent desires to achieve x, this desire will be most successfully achieved by doing y'. Even the staunchest fan of naturalized epistemology should be comfortable, I would hope, with 'normative' claims of those sorts. Ultimately, however, I don't think that these issues need to be settled for current purposes; nothing that is to come will hinge on any particular assumptions on the nature of the normative.

I also don't intend my use of terms such as 'epistemic obligation' to imply any particular kind of view on knowledge or justification. I don't mean to imply that the correct theory of knowledge is

⁶ Indeed, epistemic logic recognizes a variety of possible axiom systems, corresponding to a scale of epistemic states of different levels of demandingness.

‘deontological’, for instance. As William Alston (1988) has noted, we must take some care in characterizing epistemic normativity. If we claim that a subject is obligated/permitted/forbidden to believe such-and-so, then if we are doxastic involuntarists these obligations will appear to run afoul of the principle that ‘ought implies can’. This seems to be a general issue when one speaks of sort of epistemic obligations as placing restrictions on the sorts of *beliefs* we may hold.

To avoid this, rather than talking of obligations and permissions with regard to belief, we can instead talk of obligations and permissions with regard to things we do clearly have control over – our gathering of evidence, our use of inference procedures, our reflective consideration of our belief set, and so forth (see also Kim 1994, Feldman 2000). Call these ‘epistemic actions’. When I speak of epistemic obligation, then, I have in mind not the obligation to believe such-and-so but the obligation to perform such-and-so epistemic action. None of this requires us to take any particular position on the nature of justification or knowledge. A die-hard reliabilist, for instance, might accept that the limited reliability of human cognition frequently leaves agents with unjustified beliefs despite their having fulfilled all their epistemic obligations – in other words, despite having performed all the epistemic actions that could reasonably be expected of them.

Nonetheless, it is plausible that there is some link. If knowledge requires (say) beliefs formed by a process that is at least 80% reliable, then it might be obligatory for me to perform any epistemic action within my power until reaching that level of reliability (or something along these lines). Or perhaps my obligations are linked to my own assessment of my epistemic state – if I take myself to not yet know whether *p*, then I am obligated to investigate or reason or reflect further.⁷ A great many ways of cashing out the link between epistemic action and knowledge will leave open the possibility that I might fulfil my obligations and achieve true belief, yet fail to know; or vice versa. For the former proposal, I might perform all epistemic actions available to me, yet achieve my true belief via a merely 79% reliable process; or for the latter proposal, I might mistakenly take myself to know and hence decide that further epistemic actions are unnecessary. Again, nothing that follows will hinge on any particular take on the correct relationship between epistemic obligations and knowledge.

In any case, it seems to be epistemic actions (and not beliefs per se) that serve as the ultimate locus of epistemological praise and blame in ordinary discourse. When I criticize a friend for falsely believing that a sham medical treatment would cure her chronic illness, I would explain my disapproval by noting that she ought to have done more research before wasting her money. When I criticize another’s belief, what I really criticize is normally the fallacious reasoning that caused them to hold it, or their lack of reflection on a belief their parents instilled in them, or their cherry-picking of evidential sources, or what have you. Ultimately there must be *some* sense to be made of the familiar, ordinary talk regarding what one epistemically ‘ought’ to do, regardless of one’s philosophical positions on the various issues surrounding epistemic normativity. I don’t have any particular horse in those races; my terminology is ultimately born from convenience and expository ease. My use of various normative terms should be interpreted in whichever way one prefers to interpret the normative language used by non-philosophers in ordinary attributions of epistemic evaluation.

2.2 Epistemic supererogation and special epistemic obligations

⁷ Assuming, that is, that the truth or falsity of *p* is of interest to me.

With these preliminary caveats in place, let's consider what moral normativity might suggest to us about epistemic normativity. Ethicists standardly recognize several deontological statuses an action may have: it may be obligatory or forbidden; it may be permissible; it may be supererogatory.⁸ It is obligatory for me to refrain from murdering for pleasure; it is permissible for me to buy myself a cup of coffee; it is supererogatory for me to give up coffee and donate the money I save to charity. As we've noted, parallel categorizations feel quite natural within the realm of epistemology. It is obligatory for me to gather a certain amount of information before endorsing a presidential candidate; it is permissible for me to trust my senses; and so forth. Strangely enough, however, the category of supererogation is rarely discussed within epistemology. Tidman (1996) makes brief use of the idea; Hedberg (2014) is a rare instance of a paper-length treatment of the notion. But epistemic supererogation nonetheless remains largely unexplored territory.

We've already noted that there are many possible epistemic states that exceed the requirements for knowing. It seems equally clear that there will be epistemic *actions* that exceed our epistemic obligations – especially if our epistemic obligations are in some way tied to the requirements of knowing. Imagine again that knowledge requires one's belief to be formed by a process that is 80% reliable. If I select and employ a method which is 98% reliable, then *prima facie* this looks to be a case of epistemic supererogation – I have exceeded my epistemic duties. For a more concrete example, suppose that I check ten separate sources before resting content with my belief that Mogadishu is the capital of Somalia. It seems obvious that I did not *need* to do that; surely one or two would have been enough (indeed, in actual fact I only checked one while writing this paragraph). But it also seems obvious that I have improved my epistemic position at least a little by the additional checking - though, all things considered, my time might have been better spent in other ways. The epistemic action of decuple-checking was epistemically supererogatory.

Within ethics, philosophers often discuss circumstances under which one incurs special moral obligations for one reason or another. It's intuitive, for instance, that I have greater obligations to my immediate family than to strangers. Another quite plausible idea is that members of certain professions are obligated to perform actions that would ordinarily be merely supererogatory. A professional firefighter, for instance, is obligated to enter a burning building (at potential personal risk) to save a trapped inhabitant. Yet for the civilian passerby, such an act is merely supererogatory. Similarly, doctors are under moral obligations that do not apply to ordinary folks, such as the obligation to treat, to maintain confidentiality, and to adhere to the particular procedures surrounding informed consent. Professors, as well, are under certain special moral obligations - e.g. to their students. Plausibly, the members of these professions voluntarily take on these extra moral obligations when they join their field. Perhaps somewhat more speculatively, the very existence of these professions serves to increase instances of morally supererogatory acts within society; we pay firefighters so that we do not need to rely on the fickle benevolence of strangers.

Might there be a similar phenomenon within the realm of the epistemic? I claim that there is: professionals in many academic and inquiry-centered fields take on extra epistemic obligations in virtue of their membership in their profession. In other words, what would ordinarily be *epistemically* supererogatory becomes, for the professional, epistemically obligatory. Recall, for instance, the aforementioned stringent testing procedures that regulate the adoption of new clinical treatments; I

⁸ And of course, actions can belong in more than one of these categories.

would argue that this indicates that medical researchers are under extra epistemic obligations with regard to propositions involving the efficacy of various drugs. They are obligated to perform more epistemic actions with regard to those propositions than the rest of us are, and they incur this special obligation as a result of their professional role.

Medical testing is far from the only example. There exists a vast array of discipline-specific methodological requirements that *prima facie* suggest the existence of extra, non-standard obligations to perform epistemic actions. The most obvious examples come from the sciences. Consider for instance the measurement of significance employed in analyzing experimental data – the p-value. A scientist cannot infer directly from the raw experimental data to the support of her hypothesis; her experimental results are only evidence for the hypothesis if it can be shown that the likelihood of such results manifesting given the ‘null hypothesis’ is sufficiently low. Interestingly, the required threshold differs from field to field. In the social sciences, data is generally required to meet a significance threshold of .05, reflecting a 95% confidence level. Yet within the field of particle physics, researchers were unwilling to announce the discovery of the Higgs-Boson particle until reaching a five-sigma confidence level, corresponding to a p-value of 0.0000003.

Now, it is obvious that we do not expect anything like this analysis of significance from ordinary cognizers in everyday circumstances, even for information to which it could be applied. For one, the vast majority of cognizers have no idea how to perform such analyses. But even for those that possess the relevant skill, it would be extraordinarily strange to apply it to, say, determine whether one ought to believe that a San Antonio Spurs loss tends to cause one’s husband to be in a foul mood. It’s simply not worth the time or effort to achieve the increase in epistemic standing that would result from an organized, methodical gathering and analysis of post-game mood data. Casual observation is sufficient.

Further examples of extra-stringent requirements in science are easy to generate. The standards of experimental design, for instance, incorporate many procedures and restrictions that serve to lessen the impact of various cognitive errors and biases. Double blinding is a familiar example; it serves to lessen the biasing effect of the experimenters’ and subjects’ expectations. The practices of randomization and replication, the use of control conditions, and so on also serve to reduce the possibilities of error. Scientists also take greater care than laypersons to minimize the fallibility of perception, as evidenced by the use of precise measuring apparatus and repeated measurements, as well as e.g. the use of video recording, multiple independent coders for data, and so forth. And even after all this methodological care, scientists frequently perform meta-analyses to reduce the risk of error yet further. It goes without saying that these practices are not standard during everyday cognitive activity, even in circumstances where they could be applied – and even for those few who are trained in their use.

The sciences are the clearest example of elevated methodological requirements, but they are far from the only case of apparent special epistemic obligations. Consider the epistemic practices enshrined in our legal systems. The US federal courts, for example, make use of a complex system of rules of evidence which regulate admissible evidence in court proceedings. ‘Hearsay’, in the technical sense of a statement based on testimony received outside the courtroom, is prohibited as admissible evidence in court; it is, of course, a perfectly normal source of evidence in daily belief-formation. Similarly, evidence regarding a defendant’s character is admissible only in certain well-defined circumstances. Many readers will be familiar with the goal of demonstrating guilt ‘beyond a reasonable doubt’; in fact, this is

but one standard of evidence among many in use across various US legal contexts.⁹ Most of us, of course, tend to form the vast majority of our beliefs without meeting this rigorous standard.

I won't belabor the point, other than to note briefly a few other fields that strike me as employing elevated standards. Journalists are plausibly under special obligations regarding fairness, careful sourcing and verification of evidence, fact-checking and the like. Mathematicians may well be under an obligation to seek proof of mathematical propositions in any case where this is possible. Note, for instance, the continuing quest to prove Goldbach's conjecture, despite the fact that computers have provided inductive evidence that the conjecture holds by verifying it for all primes below 4×10^{18} .¹⁰ Academics in all fields, including philosophers, are under all sorts of obligations regarding clarity and rigor of argumentation, pursuit of relevant literature, and so forth. Consider for a moment the vastly differing standards to which you hold an undergraduate and a colleague – the very same paper that would receive an A in an introductory class would provoke scathing criticism, even reproach, when presented by a senior professor. Professor McX really *ought* to be ashamed; he *should* have been more familiar with the work of so-and-so; he *should* have considered such-and-such obvious objection.

2.3 Explaining the phenomenon

So let's suppose that I have convinced you that there is at least a *prima facie* difference between the epistemic activities that ordinary knowers standardly perform, and the stringent methodological requirements that professional inquirers hold themselves to. There are of course a number of ways in which we might explain this difference in expectations; I've mentioned shifty-standard accounts of knowledge, but another obvious route would be to claim that the extra obligations faced by the scientist are moral or practical, rather than genuinely epistemic. Finally, one might claim that everyday cognizers are, despite appearances, obligated to e.g. reach the level of statistical significance demanded in the sciences. We'll deal with all these possibilities shortly. For the moment, let's focus on a question that arises if we take the phenomenon at face value - why should this difference exist?

To my eyes, it is for essentially the same reason as in the moral case. In ethics, a common objection to naïve utilitarianism is that it is obscenely over-demanding; it would require each of us to give up on all of life's small luxuries, forgoing our daily lattes so that we might donate every scrap of non-essential income to the alleviation of world poverty. Even utilitarianism's patron saint Peter Singer falls massively short of such a demand. We are mere mortals, and the moral obligations that hold of us must respect our weaknesses. Nonetheless, in certain professional contexts we demand at least a somewhat higher ethical standard; we pay firefighters, doctors and the like to perform actions that are not morally required of the layperson.

If this is true in the moral case, then surely the same may be true in the epistemic case. Consider again the p-value threshold required of data in the sciences. Even meeting the relatively lax .05 value required in the social sciences often requires dozens of subjects; the ordinary man on the street simply

⁹ Interestingly, civil cases frequently only require a 'preponderance of evidence', which consists in the proposition in question being more likely true than false. This is quite plausibly a lower standard than required for knowing. To my eyes, this simply indicates an even greater variety of possible epistemic standings of interest.

¹⁰ There's a thorny question regarding whether inductive evidence can ever grant even ordinary knowledge of mathematical claims. I'm inclined to say that we are justified in believing Goldbach's conjecture to be true; readers who disagree may simply ignore this example.

does not have the luxurious quantities of free time that would be needed to gather that much data. If doing so were obligatory before forming beliefs about causal dependencies, most working folk would be in flagrant dereliction of duty. If the link between epistemic obligation and knowledge is reasonably tight, this would further imply that most of us know very little. Similar issues arise with the knowledge norms of assertion and action. If we refused to assert causal dependencies when bereft of scientific-grade evidential wealth, communication would be extraordinarily tricky. And if we refused to act on causal beliefs when so bereft, we'd likely go extinct. We are creatures of limited time and means. We have to make do with moderate levels of epistemic risk.

Yet the limitations that we face as individuals are much less relevant when considering the scientific community as a whole. Ordinarily, we can't expect the average man on the street to implement the stringent methodology of modern science; but we can if we pay him to spend 40-plus hours a week doing so. For certain tricky propositions it may take years, decades, or even longer to reach scientific-grade knowledge; but since science is not the task of a single individual, time constraints largely fall to the wayside. Similar points could be made about other resources that limit inquiry, such as cognitive energy, access to information and education, funds to purchase needed equipment, and so forth. We can reasonably expect quite a bit more from a group of well-trained professionals who devote much of their waking hours to inquiry than we can from a solitary individual who must balance inquiry with the other demands of daily living. By introducing professions whose charge is to conduct high-quality inquiry, we make the unreasonable reasonable - and the supererogatory obligatory.

3. Special epistemic obligations and the aims of professional inquiry

As noted earlier, I think it is plausible that the epistemic activities that we consider obligatory in ordinary cognizing are linked, in some way, to the goal of achieving knowledge. Perhaps knowledge is our ultimate goal, or perhaps this goal is merely instrumental to the ultimate goal of achieving truth. But we might say that, in ordinary inquiry, it is at least plausible that the *proximal* aim is to know, and we expect cognizers to take appropriate epistemic actions towards achieving that goal. The epistemic obligations present in inquiry-centered professions, however, appear to be different than those active in everyday contexts. This suggests that the goal may be different, too. I'd like to suggest that professional inquirers don't aim at *mere* knowledge, any more than a philanthropist aims at (say) merely avoiding causing harms.

Given the plausible importance of knowledge, there is an obvious temptation to assume that all of our activities of inquiry aim for it (if only instrumentally, as a means to get at truth). Those who take knowledge to be the aim of ordinary epistemic activity are thus likely to hold that knowledge is also the aim of science, journalism, law, philosophy and so forth; they will likely be less than happy with the proposal I have just made. Consider for instance Alexander Bird's statement about scientific inquiry: "as Aristotle tells us in the first line of the *Metaphysics*, 'All men by nature desire to know'. Science is the institutional embodiment of this desire" (Bird 2008, 281). As Moti Mizrahi (2013) notes, this knowledge-based conception of scientific progress seems to reflect the way scientists themselves talk about their work; and it was plausibly the default view among philosophers until Kuhn's *Structure of Scientific Revolutions*.

Bird's arguments for the knowledge view, however, are not aimed at the antirealist who denies scientific progress, but at the realist who conceptualizes progress solely in terms of accumulation of

truths. Bird (2007) notes that we would not consider science to have ‘progressed’ if the scientific community began to believe a truth for irrational reasons; he takes this to imply that scientific progress consists in accumulation of knowledge. Assuming that the aim of science is to achieve that which constitutes scientific progress, then if Bird is right then the aim of science is knowledge rather than true belief. One could easily imagine similar arguments for the knowledge aim in other professional fields.¹¹

There is, however, a certain failure to consider the full range of possibilities here. Why should the realist’s options be limited to truth or knowledge? Suppose we accept Bird’s argument against truth as the standard of progress; why think that scientists should be satisfied with mere *knowing*? Why shouldn’t their goal be (say) certainty, or knowing that one knows? Perhaps all men do desire to know; but don’t some of us also want more? Can’t we, like the philanthropist, aim higher than the average man on the street? Bird takes science to be the institutional embodiment of the desire for knowledge; I’d suggest it is an institution whose existence serves to promote epistemically supererogatory practices of inquiry.

Nonetheless, there are multiple ways one could explain the epistemic practices of professionals without departing from the idea that knowledge is the aim of science (and journalism, and philosophy, and so on). There seem to be two broad options. First, one might deny that the epistemic obligations of professionals genuinely differ from those of laypersons. Second, one might accept that the obligations differ but hold that these obligations can be reconciled with the view that even professional inquiry aims at knowing. Let’s examine each of these in turn.

3.1 Objection: the obligations are not epistemic

Taking the first of these aforementioned alternate explanations first, one might wonder whether the apparent special obligations described in the last section are genuinely epistemic. Perhaps they are merely pragmatic obligations, or even moral obligations. And indeed in the field of medicine the moral option has prima facie plausibility; there is presumably a moral obligation to prevent the release of harmful pharmaceuticals to the public, for instance. But most other fields of inquiry aren’t like this. There are no plausible moral obligations surrounding the study of, say, genetic inheritance patterns in insect populations in Namibia.

Are the oughts that drive the professional inquirer’s epistemic actions are merely prudential or pragmatic? Again, much of the work done by professional inquirers has no immediate practical benefit; philosophers know this better than anyone. Of course, there are special pragmatic reasons for an individual researcher to follow the then-current expectations regarding blinding and the like – if she does not, she will not succeed in publishing, which will impede her chances for funding and so forth. But researchers do not just follow *current* methodological expectations. They also alter their methodological expectations in response to new information. The introduction of placebo-controlled trials is a case in point; this alteration of methodology would make little sense if researchers were only concerned with ‘following the pack’. We should also recall that everyday inquiry frequently has quite a bit of pragmatic utility. If the obligations that surround ordinary epistemic activity are epistemic rather than pragmatic, then surely those surrounding professional inquiry are as well.

¹¹ See for instance Kelp (2014) for somewhat similar arguments in favour of the view that knowledge is the aim of inquiry generally.

3.2 Objection: the obligations are universal

A second way to deny the existence of special epistemic obligations for professionals is to claim that *all* inquirers, even laypersons, are obligated to meet the standards professional inquirers hold themselves to. This need not involve claiming that all inquirers are obligated to run double-blinded experimental studies; instead, it might simply involve the claim that all inquirers are obligated to ensure that their beliefs are free from cognitive bias and other reasoning errors. We do, after all, criticize people who form beliefs in biased ways. A biased belief, so the story would presumably go, is unjustified and thus patently not knowledge. Given the plausible link between knowledge and epistemic obligation (in ordinary contexts), it would then be a short step to the claim that ordinary inquirers are generally obligated to act so as to avoid bias.

Of course, it is true that cognitive biases and other errors are frequently incompatible with knowledge. If the gambler's fallacy leads me to believe that I will win the next round of blackjack, I do not know that I will win the next round of blackjack – even if it turns out that my belief happened to be true. But a bit of reflection will show that at least *some* cognitive flaws must clearly be compatible with knowledge, if our ordinary ascriptions of knowledge are largely on the right track. Take for instance base-rate neglect. Textbook examples of base-rate neglect typically involve medical tests which have a certain rate of false positives; in order to calculate the probability that one has disease x given a positive test result, one must make use of information about the base rate of disease x in the population. Even doctors, however, are prone to ignore this information, thereby putting too much weight on positive test results. Now the average patient has likely never given a single thought to how base rates affect the evidential value of a clinical test; yet we display absolutely no hesitation in attributing knowledge to those who have formed a (true) belief that they have such-and-so disease on the basis of such tests.

Other common biases, such as confirmation bias, bandwagon effects, overconfidence, and use of non-optimal heuristics such as availability and representativeness strike me as equally compatible with knowledge, at least in some cases. What's more, some of these 'errors' are even arguably adaptive, enabling faster decision-making with less expenditure of cognitive effort. Indeed, given the ubiquity of such cognitive behaviors, withholding knowledge attributions from any belief that had been 'tainted' by said behaviors would likely leave most cognizers with a rather impoverished knowledge base.

Finally, consider the following: obviously there will be many cases where a person has formed a belief while influenced by bias, and where that belief is therefore not justified. But for any such case we can imagine a parallel case where the bias is not present, but where the agent has performed no particular epistemic actions to *prevent* the operation of bias, either. For instance, imagine a case where an agent's judgment of her colleague's competence is not in fact influenced by the colleague's ethnic background, but where the agent has not made any particular effort to shield herself from possible tacit racial bias. I'd argue that cases of this sort are compatible with justified belief. Most people *are* generally in this latter position – very few people are even aware of the existence of cognitive biases, much less actively working to prevent their influence. Yet note that in fields of professional inquiry, it is not just obligatory to *be* free of bias; it's obligatory make an *active effort* to *ensure* one is free of bias. This obligation simply does not hold (at least to the same degree) in ordinary cognitive activity.

3.3 Objection: the obligations reflect an attempt to maximize knowledge

This might lead to another thought: perhaps scientists and other such inquirers are aiming at knowledge, but are simply going to extra lengths to make absolutely *certain* that they have in fact met their aim? This might be a version of the second strategy mentioned above – we might accept the existence of elevated obligations but maintain that the aim of all inquiry is still knowing. But notice that, if this hypothesis is correct, professional inquirers *would* be aiming at a higher state than knowing – they would be aiming at ‘being certain that one knows’.¹²

More plausibly, perhaps professional inquirers are simply organizing their epistemic activities such that they maximize the amount of knowledge they produce. Surely, by reducing the chances of error professional inquirers improve their chances of achieving knowledge. But it is not at all clear to me that the stringent error-reducing procedures in place in these fields actually succeed in *maximizing* knowledge. Consider two hypothetical inquirers, Alfred and Betty. Both desire to maximize their knowledge. Alfred proceeds by devoting extraordinary amounts of time and energy to eliminating every possible risk of error for each of the propositions he considers. Since this takes so much time, Alfred only investigates three propositions over the course of a year. Betty is also quite scrupulous about avoiding error, but not nearly so much so as Alfred. She permits a reasonable risk of error for each proposition that she investigates, and as a consequence she investigates hundreds of propositions over the course of the same year. Chances are quite good that Betty will end up with more knowledge than Alfred at the end of the year. (Alfred, by contrast, will end up with a lower quantity of higher-grade epistemic states). Given our limited resources, there is a tradeoff between quality and quantity of epistemic goods over any given period of time. Scientists and other professional inquirers do of course allow for some degree of error, but it is in no way obvious that they are Bettys rather than Alfreds.

Alternatively, but relatedly, it might be suggested that professional inquirers employ more stringent methods solely because they are inquiring into trickier matters. Certain propositions are, after all, harder to come to know than others – it is more difficult to know that asbestos causes cancer than it is to know that the cat is on the mat (leaving cases of testimony aside, that is). Indeed, a great many propositions may be unknowable without use of the particular methods of science – from such obvious cases as ‘electrons have negative charge’ to less-obvious cases involving, say, small effect sizes in experimental contexts. Thus, an objector might claim that professional inquirers always aim at (mere) knowledge, but that the difficulty of their subject matter necessitates the use of the stringent bias and error reduction procedures we have discussed.

In fact, I do think that the existence of differing levels of ‘epistemic demandingness’ – that is, the level of difficulty inherent in coming to know a proposition – partially explains the existence of special obligations within contexts of professional inquiry. Nonetheless, this cannot be the whole story, for there are many propositions which are both knowable by everyday means and subject to the special elevated standards we’ve been discussing. A substantive proportion of the propositions studied by medicine and by psychology fall in to this category; particularly those bits of ‘common sense’ that are later verified scientifically, such as e.g. the claim that fatty foods contribute to obesity or the claim that financial hardship causes psychological stress. It goes without saying that philosophy, as well, vigorously

¹² And presumably, being certain that one knows that p entails being certain that p. Which, as we have agreed, is a much more demanding state than mere knowledge that p.

and carefully investigates a large number of claims that fall within the domain of common sense knowledge.

So why do professionals bother investigating these sorts of common sense claims with the full rigor of their fields' methods, if they are already known? In part, because of the very large number of 'common sense' claims that are later shown to be false. Examples here from the sciences include the claim that cold weather causes colds, or that 'letting it out' is a beneficial way to cope with feelings of anger. These bits of common sense are false, but that should not lead us to claim that all common sense fails to be knowledge; again, that would be too skeptical. A better conclusion is simply that we do not always know which bits of common sense are known – in line with the quite plausible claim that knowing does not entail knowing that one knows. In order to separate the wheat from the chaff, it is worth the professional's time to double-check even propositions that might seem obvious. On this perspective, professional inquirers may well be aiming at knowledge of knowledge rather than 'mere' knowledge.

But note that we need not be forced into shoehorning all goals of professional inquiry into iterations of knowledge. Perhaps certain professions aim to be certain that a proposition is known, or to be 99% confident that it is known, or even simply to be more justified than would be required in ordinary contexts. For whatever reason, propositions that are plausibly known prior to professional inquiry are frequently investigated in a more rigorous manner by professionals. This suggests that the special obligations of professional inquirers are not directly determined by the epistemic demandingness of the investigated propositions.

3.4 Objection: the obligations reflect shifting standards for knowing

Let's turn to the second strategy mentioned above – that of admitting obligations while maintaining that the aim of all inquiry is knowledge. The most obvious option here is to invoke a shifting-standards account of knowledge. Let's begin with contextualism. Contextualist accounts of knowledge claim that knowledge attributions – that is, statements of the form "S knows that P" – express different propositions depending upon the context in which they are uttered. In this way, terms like 'knowledge' and 'knows' are said to resemble other familiar context-sensitive terms such as 'here', 'big', and so on. What determines the meaning of 'knows' in a given utterance is the features of the attributor's conversational context. If a skeptical scenario has recently been mentioned, for instance, then 'knowledge' comes to express a more demanding epistemic property than it would in a more everyday context. One might think that the elevated standards operative in the fields we've been discussing could be accounted for quite easily by a contextualist theory of knowledge. The story would go something like this: during conversations in these fields, the possibility for error is typically made salient, thus resulting in a shift to a more demanding standard, which then generates special epistemic obligations.

There are a few difficulties in taking this route, however. First, there are certain sentences which should turn out infelicitous on this sort of account which nonetheless strike me as fairly natural. Suppose, for instance, a scientist discussing possible future research projects with a colleague says something like the following: "We've known for a long time that long work hours cause mental fatigue in employees, but we don't yet have scientific evidence demonstrating this effect". On a contextualist

account, the scientific context of this conversation should raise the standard for knowing such that the scientist's knowledge claim is inapt. Yet it seems perfectly normal to claim, in a single breath, that a piece of 'everyday' knowledge has not yet been subjected to rigorous investigation that meets the standards of scientific inquiry.

Second, it's not clear that there is a *single* standard in effect in a given scientific context in the first place. Within, say, a psychology laboratory, the investigation of certain propositions will be subject to extra-stringent methodological requirements; but other propositions will be held to essentially the same standards that are operative in everyday contexts. The psychologist will not feel obligated to perform any extra epistemic actions before, say, making a claim about the number of subjects in her control group. This complexity is difficult to accommodate on a contextualist approach.

A further difficulty seems to arise from the fact that, on a contextualist account, it is the conversational context of an *attributor* that determines the standards. Thus, suppose researcher Smith has performed a fairly shoddy experiment to test the effectiveness of drug X; she has neglected to employ a placebo, and she did not properly randomize the assignment of subjects to experimental and control groups. Nonetheless, let's suppose that her research provides better evidence than mere casual observation of the drug's effectiveness during clinical practice. If we try to extend the contextualist account to cover obligations to perform epistemic actions, it seems that should be conversational contexts in which it would be true to claim that Smith has done nothing epistemically blameworthy (conversations by disinterested laypersons, perhaps). But it seems quite clear that any such claims would be false, no matter the conversational context. She does deserve blame; she has neglected her duties as a researcher.

Ultimately, the contextualist solution is problematic because it is at heart a semantic thesis rather than a thesis about epistemic norms. Contextualism tells us about the meaning of the word 'know'; it does not obviously tell us how we ought to go about forming beliefs. On a contextualist account, any change in standards is an artefact of conversation; but not so for the scientific standards we've been discussing. The epistemic obligations of the scientist are what they are. One can't change them through conversational maneuvering.¹³

A more promising approach would be to appeal to features of the inquirers themselves, rather than any particular knowledge-attributor, to account for the shifting standards. The main competitor to contextualism within multiple-standards accounts of knowledge, which is standardly known as either Subject-Sensitive Invariantism or Interest-Relative Invariantism, therefore suggests itself as a plausible explanation for the phenomena at hand. According to SSI accounts (as I'll call them), the meaning of 'knows' does not change in different contexts. Instead, whether an agent knows is simply sensitive to non-epistemic features of the agent – typically, what is at stake for the agent. If it is terrifically important that the agent have a true belief regarding proposition p, then the level of justification the agent must have before knowing p is elevated.

¹³ I have here largely assumed a version of contextualism upon which the standards are determined by conversational salience. This is not the only possible contextualist approach, of course – however, insofar as contextualist theories hold that the truth conditions of knowledge attributions are in part determined by the attributor, this last argument stands. The norms of inquiry simply do not appear to be attributor-sensitive.

Oddly, proponents of SSI accounts have tended to argue for them on the basis of semantic considerations, just as proponents of contextualism often do. From my perspective, however, whether or not SSI accounts properly capture our usage of the English word 'knows' is of very little interest. Of much greater interest is the overwhelmingly plausible observation that one's epistemic obligations are affected by what is at stake. If your life depends on depositing a check on Saturday, you *ought* to double check the bank times. An agent who fails to do so in such a case deserves quite a bit of blame; an agent who rests content with her memory of previous bank visits while under no particular practical pressures seems, by contrast, blameless. Regardless of whether we agree that 'knows' picks out a property whose presence is sensitive to the practical stakes of a subject, everyone should admit that what epistemic actions one *should* perform is sensitive in such a way.¹⁴

Supposing that there is at least a moderately close link between epistemic obligation and the nature of knowledge, the sensitivity of epistemic obligations to stakes naturally suggests the sensitivity of knowledge to stakes. But what of the unconventional epistemic obligations of professional inquirers? Here, stakes-based SSI accounts don't seem to be quite general enough to capture the phenomena. One obvious difficulty is that scientific curiosity and high practical stakes don't very closely correlate. Medical research obviously involves high stakes due to the risk of harming a patient population; but the vast majority of scientific research is not accompanied by such immediate and obvious practical consequences. If physicists are wrong about the Higgs-boson, very few people will be harmed.

There may, of course, be very high stakes for an individual researcher with regard to a given scientific proposition; her professional career might rest on getting things right with regard to that proposition, particularly if she is not yet tenured. But SSI gets things wrong here, too – the methodological standards of scientific inquiry aren't higher for the fresh assistant professor than they are for the jaded emeritus researcher. As Jonathan Schaffer has noted, "One cannot gain a competitive advantage in scientific inquiry... by not caring about the result" (Schaffer 2006, 96).

Nonetheless, perhaps one could imagine *some* form of shifting-standards account of knowledge that would mesh comfortably with the apparently elevated epistemic obligations of professional inquirers. Perhaps even some version of a deontological view of knowledge, to the effect that one knows *p* when and only when one has a true belief that *p* and in addition one has fulfilled all one's epistemic obligations with regard to *p*. On such an account, the standards for knowing would be elevated when one has special epistemic obligations.

However, even leaving aside the issues surrounding doxastic voluntarism mentioned earlier, this way of accounting for the phenomena has a particularly undesirable consequence. It implies that there are a number of propositions known by laypeople but not by specialists, despite the specialists having done a much more thorough job investigating such propositions. A layman with anecdotal evidence that such-and-so herb is effective against colds might possess knowledge, while a scientist mid-way through a meta-analysis of the herb's effectiveness might lack said knowledge, despite possessing a belief based

¹⁴ One could, however, argue that the 'should' here has merely practical and not specifically epistemic force. Nonetheless it seems to me uncontroversial that one is obligated, for one reason or another, to perform certain extra epistemic actions under circumstances of increased practical stakes. In any case, if the 'should' in high-stakes cases is merely practical, then presumably it is merely practical in low-stakes cases as well (else, we face thorny question of exactly *how high* the stakes must be before our obligations become practical rather than epistemic); this, then, would suggest that all epistemic obligations are merely practical.

on much greater evidence. Indeed, it seems every shifty-standard account of knowledge will be forced to accept this sort of conclusion – even versions of contextualism and SSI that might avoid the other objections raised above.

But suppose we were to bite that particular bullet. Ultimately, I am arguing for a ‘higher-than-knowledge’ aim because I believe that there are epistemic categories beyond knowledge (and the other ‘usual suspects’ like true belief) which are worthy of philosophical attention. This, I think, is correct regardless of debates over aim. Even if one manages to defend the claim that professional inquirers aim at (mere) knowledge, it will remain true that knowledge is not the most theoretically interesting epistemic category when it comes to epistemological questions surrounding these professions. Thus, it will remain true that epistemologists should concern themselves with epistemic categories beyond those that are currently standard within traditional epistemology.

Here’s why. Note first that even if a shifting-standards account of knowledge is maintained, we can still speak of ‘the epistemic state one has when one meets the standards for knowing active in scientific contexts’ (e.g.). Call this knowledge_s. Knowledge_s is a type of knowledge, just as red is a type of color. And just as there might be circumstances where one is interested in the nature of red rather than in the nature of color generally, there might easily be circumstances where it is knowledge_s, rather than knowledge generally, that is of interest. Similarly for the subtypes of knowledge that correspond to other areas of professional inquiry.

As an example dear to my heart: suppose one is concerned about the use of intuitions in philosophy. One *might* ask whether intuitions are reliable enough to produce knowledge. But it’s entirely possible that intuitions are sufficiently reliable to generate knowledge in everyday contexts, while being insufficiently reliable to meet the elevated standards active in philosophical contexts. What one should *really* be interested in is whether intuitions are sufficiently reliable to generate the subcategory of knowledge that philosophical contexts demand – ‘knowledge_p’, say. Defending the use of intuitions in philosophy by defending their status as a source of knowledge is rather like arguing against the necessity of telescopes in astronomy by defending naked-eye perception’s status as a source of knowledge. Sure, unaided perception is a sufficiently high-quality epistemic source to generate knowledge – but in certain contexts of inquiry, it is just not good enough.

4. Conclusion

I’ll end with a few final clarifications and elaborations on the view I’ve outlined. Broadly, the view I’ve defended is that professional inquirers face more stringent epistemic obligations than layfolk; I’ve argued, further, that this strongly suggests that they aim at a higher epistemic status than mere knowing. I don’t, however, have much of a position on whether or not this higher aim is merely in service of an ultimate goal of achieving true belief. It’s possible that the higher epistemic standard is merely instrumental. I’m satisfied if I merely convince you that the *proximal* goal of scientific activity exceeds the requirements for knowing. In fact, as just mentioned, I’m satisfied if I haven’t even done that – so long I’ve convinced you that, for some epistemological purposes, epistemic states that exceed knowing may be the appropriate target of philosophical investigation, and that contemporary epistemology’s focus on knowing is therefore problematic.

Some readers might wonder whether the goal of professional inquiry might be something like understanding. I have no problem with the idea that understanding is a valuable epistemic state;

however, I don't think that the arguments I've given in this paper support the idea that understanding is the aim of these vocations. The methodological obligations I've discussed aim at reducing the chance of error; but presumably one could have even total immunity to error while lacking understanding. What's more, on the current picture talk of 'the' goal of professional inquiry is rather misleading. As we've seen, there are different epistemic expectations corresponding to different fields and even to different questions or projects within a field. This implies that the epistemic states that satisfy e.g. physicists are different from those that satisfy e.g. journalists, and that the epistemic goal of a physicist may vary depending upon the particular project upon which she is embarked or the particular proposition she is considering.

The picture all this suggests is much more pluralistic – vastly more pluralistic than is usual in 'traditional' epistemology. There are exceptions – Stephen Hetherington (2001), for instance, has advocated viewing knowledge as a spectrum, thereby recognizing that one might know one proposition 'better' than one knows another – in other words, recognizing a variety of possible epistemic standings. But largely, traditional epistemologists have had a puzzling tendency to neglect the vast majority of possible epistemic states in favor of a few supposedly crucial distinctions. Contrast traditional epistemology's 'on-or-off' approach to belief and knowledge with e.g. a Bayesian framework's spectrum of credences; the latter type of approach fits much more comfortably with the complex, diverse array of norms and obligations that appear to hold in fields of professional inquiry. In fact formal epistemology, quite generally, seems much more amenable to the sort of pluralistic picture I have in mind; my complaint in this paper is restricted to the over-emphasis on knowing present in traditional epistemological projects.

A few final comments in a somewhat more speculative vein. If the picture I've defended is roughly correct, one consequence is that the correct characterization of knowledge simply cannot be got at *a priori*. 'Knowledge', I've urged, must pick out an epistemic threshold that we mere mortals can fairly regularly attain in our everyday cognitive lives. Else, the epistemic norms that are associated with knowledge would be uselessly overdemanding. But this threshold is obviously dependent on contingent, empirical facts about our psychological capacities as a species, about the environment we find ourselves in, and so forth. We cannot know what knowledge requires without knowing a great deal about, for instance, the various shortcomings of our perceptual capacities.

What's more, depending on one's semantic views, the fact that knowledge is of any interest at all to us may turn out to be utterly contingent. It seems plausible that if our cognitive abilities had been different, we would have been interested in a different epistemic standing. For instance, if we had possessed god-like intelligence and phenomenally long lifespans, perhaps we might have 'hit upon' a state closer to certainty as our core epistemic concept, rather than the more modest sort of state that 'knowledge' picks out in the actual world. By contrast, if a single act of seeing took hours rather than a mere instant, we'd likely to be forced to form beliefs on less evidence; in such a scenario our core epistemic concept would presumably pick out a state less demanding than knowledge. Thus, if 'knowledge' rigidly designates, then it seems knowledge is only contingently of interest.¹⁵ This is not to

¹⁵ A reviewer notes that this claim at least *prima facie* sits poorly with views like Williamson's, according to which knowledge is of special interest due to its status at the most general factive mental state; it is the most 'general' in the sense that the possession of any factive mental state entails the possession of knowledge. A full response is beyond the scope of this paper, but a rough suggestion for a reply is this: perhaps knowledge is only the most

say that knowledge is of *no* interest. What one epistemically ought to do is, I claim, constrained by contingent facts about one's abilities and environment. What most of us epistemically ought to do, in most contexts, may well be to aim to know – but in certain contexts, such as that of professional inquiry, some of us want, and plausibly ought to strive for, more.

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).

Works Cited

- Alston, W. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, 2, 257–299.
- Beecher, H. K. (1955). The powerful placebo. *Journal of the American Medical Association*, 159(17), 1602–1606.
- Bhatt, A. (2010). Evolution of clinical research: a history before and beyond James Lind. *Perspectives in clinical research*, 1(1), 6.
- Bird, A. (2007). What is scientific progress? *Noûs*, 41(1), 64–89.
- Bird, A. (2008). Scientific progress as accumulation of knowledge: a reply to Rowbottom. *Studies in History and Philosophy of Science Part A*, 39(2), 279–281.
- Craig, E. (1991). *Knowledge and the state of nature: An essay in conceptual synthesis*. Oxford: Clarendon Press.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3), 667–695.

general factive mental state *for creatures like us*. Creatures with vastly less reliable perceptual abilities might not be able to 'see that p' – the closest factive mental state that they might achieve might be one for which we have no concept, and which might not entail knowledge, but instead a factive epistemic state weaker than knowing (in the sense of involving a lower level of justification).

- Goldman, A., & Olsson, E. (2009). Reliabilism and the Value of Knowledge. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 19–41). Oxford: Oxford University Press.
- Greco, J. (2003). Knowledge as Credit for True Belief. In M. DePaul and L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008). Knowledge and Action. *Journal of Philosophy*, 105, 571-590.
- Hedberg, T. (2014). Epistemic supererogation and its implications. *Synthese*, 191(15), 3621–3637.
- Hetherington, S. (2001). *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. Oxford: Clarendon Press.
- Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51(3), 227–232.
- Kim, K. (1994). The deontological conception of epistemic justification and doxastic voluntarism. *Analysis*, 54(4), 282–284.
- Kuhn, T. (1996). *The Structure of Scientific Revolutions* (3 ed.). Chicago: The University of Chicago Press.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- Mizrahi, M. (2013). What is scientific progress? Lessons from scientific practice. *Journal for General Philosophy of Science*, 44(2), 375–390.
- Peacocke, C. (1999). *Being Known*. Oxford: Clarendon Press.
- Pritchard, D. (2007). Recent work on epistemic value. *American Philosophical Quarterly*, 44(2), 85–110.
- Ross, O. B. (1951). Use of controls in medical research. *Journal of the American Medical Association*, 145(2), 72-75.
- Schaffer, J. (2006). The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies*, 127(1), 87–107.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.

Tidman, P. (1996). Critical reflection: An alleged epistemic duty. *Analysis*, 56(4), 268–276.

Velleman, J. D. (2000). On the aim of belief. In D. Velleman (Ed.), *The possibility of practical reason* (pp. 244–281). New York, NY: Oxford University Press.

Wedgwood, R. (2002). The Aim of Belief. *Philosophical Perspectives*, 16, 267–297.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Zagzebski, L. (2003). The Search for the Source of Epistemic Good. *Metaphilosophy*, 34, 12–28.