

Knowledge and normality

Joachim Horvath & Jennifer Nado

Synthese

An International Journal for
Epistemology, Methodology and
Philosophy of Science

ISSN 0039-7857

Synthese


DOI 10.1007/s11229-020-02823-9



Your article is protected by copyright and all rights are held exclusively by Springer Nature B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".



Knowledge and normality

Joachim Horvath² · Jennifer Nado¹ 

Received: 17 March 2020 / Accepted: 3 August 2020
© Springer Nature B.V. 2020

Abstract

In this paper, we propose a general constraint on theories of knowledge that we call ‘normalism’. Normalism is a view about the epistemic threshold that separates knowledge from mere true belief; its basic claim is that one knows only if one has at least a normal amount of epistemic support for one’s belief. We argue that something like normalism is required to do full justice to the normative role of knowledge in many key everyday practices, such as assertion, inquiry, and testimony. The view of normality we employ to flesh out this claim is inspired by experimental work on the folk notion of normality, which suggests that folk judgments of what is ‘normal’ are based upon both statistical averages as well as normative ideals within the relevant target domain. Adopting this notion of normality to set the threshold for knowledge results in a view upon which knowledge is routinely available on an everyday basis without being a merely trivial achievement. We explore several interesting consequences of this view, including the implication that the threshold for knowing may change as, e.g., the ease of availability of information in an epistemic community changes over time. The result is a ‘shifty’ view of knowledge which nonetheless retains more stability than standard contextualist or pragmatic encroachment approaches.

Keywords Knowledge · Conceptual engineering · Epistemology · Normality

1 Introduction

Most contemporary epistemologists agree that we know lots of things: we know our own names, we know where we live and go to work, we know that grass is green, that snow is white, and that tomatoes are red. Most of the things we know are quite mundane everyday propositions, as in the examples just given. Most would agree, then, that knowledge isn’t a spectacular kind of intellectual achievement, but rather an

✉ Jennifer Nado
jennifernado@gmail.com

¹ University of Hong Kong, Hong Kong, Hong Kong

² Ruhr University Bochum, Bochum, Germany

everyday cognitive commodity that provides the informational underpinnings of most of what we do and aim for in our daily lives. Any acceptable theory of knowledge should thus imply that knowledge is ‘within cognitive reach’ and available to us on a regular, everyday basis. Let us call this general perspective the *pedestrian* approach to knowledge.

The vast majority of mainstream epistemology since the 20th century assumes the pedestrian approach. But there is an alternative way to think about knowledge, often associated with Descartes or the skeptics, that makes knowledge extremely difficult or even impossible to achieve in normal human life (see, e.g., Descartes 2017; Unger 1975). In order to know a proposition P, on this understanding, one must satisfy some extremely high epistemic standard, such as being unable to doubt that P, being certain that P is true, or being infallible vis-à-vis P’s truth. Moreover, there is even some psychological evidence that our ordinary folk conception of knowledge may be of the very demanding, infallibilist kind (Nichols and Pinillos 2018). Let us call this the *olympian* approach to knowledge.

The olympian approach is rare in contemporary epistemology; so much so that the pedestrian approach can seem downright platitudinous. But the corresponding theoretical goal of keeping knowledge ‘within cognitive reach’ is, we think, less straightforward than it might seem. Consider that quintessential expression of the pedestrian perspective, fallibilism. Fallibilists claim, against the Cartesian olympians, that infallibility is not a necessary condition on knowledge. But fallibilism alone does not guarantee that knowledge will be pedestrian; a theory might be fallibilist while still espousing standards that are too demanding to be regularly met by the proverbial man on the Clapham omnibus. A theory might demand levels of reliability that are rarely attained by ordinary folk; it might demand a level of awareness of one’s evidence that ordinary folk simply do not regularly possess. In practice, theories often sidestep the issue by leaving certain quantitative details unspecified—you’re unlikely to find a reliabilist who will tell you the exact degree of reliability that suffices for knowing. In other words, while most contemporary theories are compatible with a pedestrian approach, few directly imply it.

We take the pedestrian approach very seriously; we view it as constituting a general, high-level constraint on any acceptable account of knowledge. We hold this to be true both *descriptively* and *prescriptively*, in the following sense. Most epistemologists take their projects to be descriptive—they aim to *discover* the nature or correct analysis of knowledge. We claim that any descriptively adequate theory should constitutively imply that knowledge is pedestrian, even if this may require an error-theoretic explanation of certain olympian folk judgments. An alternative approach to epistemology is *prescriptive*. A prescriptive epistemologist would recommend revisions to our ‘folk’ concept of knowledge, in the manner of recently popular ‘conceptual engineering’ (see, e.g., Fassio and McKenna 2015). We also claim that any prescriptively adequate re-engineering of our folk concept of knowledge should imply that knowledge is pedestrian.

In this paper, we aim to sharpen this general ‘pedestrian constraint’ into a specific theoretical view about the epistemic threshold that suffices for knowing—a view that we call *normalism*. Normalism is not itself a theory of knowledge, but a framework that may be ‘filled in’ with any of a variety of views on, e.g., the nature of justification;

the normalist framework will guarantee that the resulting theory constitutively implies that knowledge is pedestrian.

In brief, the normalist proposal is as follows: the degree of epistemic support required for knowledge is a *normal degree* of epistemic support. If this strikes you as worryingly trivial, bear with us. ‘Normal’ here is intended in a very specific sense, as reflecting a particular type of classification that arguably permeates everyday evaluative judgment. It is not synonymous with ‘average’—though as we’ll see, it is tied to descriptive averages in a way that naturally suits it to the pedestrian approach to knowledge. Ultimately, we aim to convince you that normalism provides a promising framework for both descriptive and prescriptive theories of knowledge—and thus, that the notion of normality has deep relevance to epistemology regardless of one’s preferred methodological orientation.

2 The everyday roles of knowledge

Before we present the specifics of the normalist view, we’ll start by discussing why we hold the pedestrian constraint to be so central to any successful theory of knowledge. We’ll first briefly review some ways in which knowledge serves as an epistemic standard that helps us to normatively regulate many aspects of ordinary human life and everyday human cooperation. We have in mind here, for instance, the various widely (though not universally) accepted knowledge norms—of assertion, practical reasoning, and so forth. We’ll then argue that, in order for knowledge to effectively serve these regulatory roles, it *must* be pedestrian.

To begin with perhaps the most familiar example of knowledge’s regulatory roles, knowledge is often held to be a normative standard for common *speech acts*, most typically *assertion*. According to a straightforward formulation of the knowledge-norm of assertion, one should only assert that P if one knows that P (cf. Williamson 2000). Therefore, someone who asserts that P in the absence of knowledge that P would act in a way that is, in some sense, normatively objectionable or criticizable. On Williamson’s version of this proposal, the knowledge norm is moreover *constitutive* of assertion, in the sense in which the rules of chess are constitutive of the game of chess. A game where the rook is permitted to move diagonally, for instance, simply isn’t chess but instead some other game.

As a second example, it has been argued that knowledge is the normative basis for *action* and *practical deliberation*, such that one is subject to normative criticism if one acts or deliberates on the basis of beliefs that fall short of knowledge. Hawthorne and Stanley formulate this norm as follows: “Treat the proposition that p as a reason for acting only if you know that p” (Hawthorne and Stanley 2008, p. 577). Proponents of this view also tend to take knowledge to be sensitive to practical stakes (see, e.g., Fantl and McGrath 2002; Hawthorne 2004; Stanley 2005); as the negative consequences for error rise, so does the threshold for knowing that P—and for permissibly treating P as a reason for acting.

Third, knowledge is arguably the *aim of inquiry* (see, e.g., Kappel 2010; Kelp 2014; Rysiew 2012; Williamson 2000), at least of an everyday kind.¹ For example, if you inquire into when your train departs, you may permissibly terminate your inquiry once you know that, e.g., your train departs at noon. If you stop inquiring before you actually know that your train departs at noon, you would—barring unusual circumstances—be subject to normative criticism. As this example makes clear, the notion of knowledge as the aim of everyday inquiry is plausibly linked to knowledge's apparent role in practical reasoning.

Fourth, it also seems plausible that we should only rely on the *testimony* of others concerning P if they know (see, e.g., Audi 1997; Burge 1993), or at least are in a position to know (Horvath 2008), whether P is true. Edward Craig (1990) has even argued that the primary purpose of having a concept of knowledge is to enable identification of suitable informants, i.e., appropriate sources of testimony. Similarly, David Henderson (2009) claims that a central role for the concept of knowledge is to act as an epistemic 'gate-keeper' for a community, flagging particular persons as approved epistemic sources.

More generally, knowledge seems to be intimately related to issues of epistemic praise and blame. We praise those who know; we blame those who fail to know. But beyond its role in *epistemic* praise and blame, knowledge also seems an important epistemic standard for everyday attributions of *moral responsibility* and the related *moral praise* and *blame*. For example, if a subject S presses a button that will detonate a bomb resulting in the death of many people, it seems critical to ask whether S knew that pressing the button would detonate the bomb. Arguably, S would, *ceteris paribus*, only be fully responsible and thus subject to the maximal amount of blame in this situation if S knew that her action would detonate the bomb.

Knowledge, then, seems to be implicated in a wide range of evaluative practices regulating essentially everything we say and do. Not all of these everyday normative roles for knowledge may be equally plausible or uncontentious. But taken together, they do make a strong case that knowledge is an important normative standard for many everyday practices. Now here is the key point: if these everyday practices are to be well-functioning, we claim, the relevant norms must be *reasonably fair*. That is to say, ordinary human beings must be able to comply with them without extraordinary sacrifice of time, effort, and so forth. And *that* is to say that the standard invoked by these norms must be pedestrian.

By means of contrast, consider the olympian view that one only knows that P if one can exclude all possibilities of error concerning P. Such an olympian view would render everyday perceptual knowledge like 'there's a barn' or 'there's a zebra' well-nigh impossible, because ordinary people in ordinary circumstances cannot reasonably exclude certain error-possibilities, such as fake barns, zebra-disguised mules, or Cartesian evil demons—unless they have an argument of the sophistication of Descartes' *Meditations* ready at hand (and arguably not even then!).

On an olympian standard of this kind, most of our beliefs would thus fail to qualify as knowledge, and consequently most of our everyday assertions, actions, inquiries

¹ For many forms of academic or scientific inquiry, the aim or standard may in fact be more than knowledge (cf. Nado 2017, 2019).

etc. would be norm-violating. Thus, it would seem, proper participation in our relevant evaluative practices would involve a willingness to criticize nearly *every* human assertion, action, inquiry, etc. This is clearly descriptively inadequate; though we do indeed criticize people for asserting, acting, or inquiring when their epistemic position is too low, we don't actually find most everyday assertions, actions, or inquiries criticizable.

Of course, all reasonable proponents of knowledge-norms claim that breaking said norms can be, in certain circumstances, all-things-considered acceptable (see, e.g., Williamson forthcoming). To take a moral example, a father who steals medicine for his child is still violating the norm against stealing, even if his action is all-things-considered morally acceptable. So, perhaps some conflict of norms might explain why we in practice exonerate most assertions as permissible. But this strikes us as implausible. Uncontroversial cases of this sort of normative 'overriding' are *exceptions*; we should not expect a genuine, practically important norm to be overridden in the majority of the cases to which it applies.

Olympian standards of assertion are not merely descriptively inadequate; they are also prescriptively inadequate. Our assertion practices, we claim, would not be improved by deeming most assertions criticizable—especially if avoiding criticism is out of the ordinary subject's reach. What would be the point of such a stringent, high-unreachable norm? There is little practical value to, say, judging secondary-school writing assignments by the standards used to select Nobel laureates in literature; *mutatis mutandis* for an olympian norm of assertion.

One might object here by appealing to the moral domain to motivate the possibility, or even the utility, of norms that most people fail to uphold. For instance, one does not refute the claim that eating meat is wrong by pointing out that the vast majority of humans eat meat. But the two cases strike us as disanalogous. There's a plausible case to be made that the world would be a better place if everyone stopped eating meat; by contrast, it seems deeply implausible that our assertion and action practices would be improved if people refrained from asserting or acting on the vast majority of their beliefs.

The above considerations apply equally well to knowledge's other normative roles, and already give us good reason to think that well-functioning knowledge-norms will be incompatible with an olympian approach to knowledge. But things are arguably worse if we view knowledge-norms as *constitutive* of their respective practices, as, e.g., Williamson does with assertion. To continue the chess analogy for constitutive norms: if nearly everyone moves their rooks diagonally (and always has), why should we think that the game they are playing is constitutively regulated by a rule that specifies horizontal/vertical rook movement? Either chess is not so regulated, or everyone has been playing a different game all along. Similarly, if knowledge is quite rare, then either knowledge is not the constitutive norm of assertion,² or human communication is in fact characterized by a *different* practice, say assertion*.

The moral of all this is that the pedestrian approach isn't just a corollary to the intuitiveness of the sorts of mundane knowledge attributions with which we opened

² Could a fan of the olympian approach simply deny that knowledge plays these normative roles in our lives? She could, but we think that this would take away much of the theoretical interest that knowledge invites. In other words, we hold that these everyday normative roles form a large part of the reason for being interested in knowledge and having a knowledge-concept in the first place.

this paper. The pedestrian nature of knowledge is crucial to its ability to fulfil the normative functions it plays in everyday life. Knowledge's being pedestrian is not merely an interesting happenstance feature; it is arguably as central to knowledge's nature as is, e.g., its factivity. As such, we claim, it deserves to be explicitly incorporated into our theories of knowledge.³

3 Normalism about knowledge

We've argued that in order for our concept of knowledge to effectively play the various roles it plays in our evaluative practices, the standards for knowing must not be overdemanding. But how can we ensure that we don't set our standards too high? The only way to successfully do so is by looking to the actual, contingent features of our epistemic abilities and resources. In other words, by looking at what is epistemically *normal* for beings like us.

Various aspects of this basic line of thinking—that a well-functioning knowledge concept must avoid an overdemanding threshold for knowing, and that to so avoid requires attending to what is normal for beings like us—have already occasionally surfaced in the epistemological literature. We are not, for instance, the first to suggest that a pedestrian standard is mandated by the practical functions of the concept of knowledge:

... we are held back from skepticism because demanding infallibility from our informants would undermine the practical requirements that comprise the purpose of our concept of knowledge. ... By setting the standard too high we would frustrate our communal epistemic practices, which is antithetical to the point of epistemic evaluation. (Hannon 2019, p. 47)

Nor are we the first to suggest tying the threshold for knowing to what is normal:

We are social animals. One's linguistic and conceptual repertoire is heavily influenced by one's society. The society will tend to adopt concepts useful to it. A concept of epistemic justification that measures the pertinent virtues or faculties of the subject relative to the normal for the community will be useful to

³ A few authors have championed the claim that olympian approaches to knowledge are consistent with the role of knowledge in the sorts of practices we review here. Davis (2007) and BonJour (2010) both suggest that our knowledge attributions and our tendency to deem most assertions and actions acceptable might reflect a 'loose' or 'approximate' usage of 'know'. Fassio (2018) similarly suggests that knowledge attributions may frequently be strictly speaking false, but warrantably assertable. However, we're still left with a pressing need to delineate the threshold at which a knowledge attribution or an acceptance of assertion/action is 'close enough'. In other words, what is the threshold for 'approximate knowing'? We hold that a pedestrian threshold is still needed in order to enable functional practices of assertion, action and so forth. Imagine a student who is told that only a perfect score in a course suffices for a passing grade, but that the professor will also grant course credit to those who are 'close enough'. Then, the crucial information is clearly what counts as 'close enough'. An ecumenical suggestion would thus be that readers committed to the views above may amend our claim in this section to '*approximate* knowledge must be pedestrian', and *approximate* knowledge is what is really of philosophical interest in understanding everyday knowledge attribution and normative epistemic evaluation'. Yet, if most of the philosophical heavy-lifting is done by approximate knowledge anyway, this also raises the question what—on balance—really favors the olympian over the pedestrian approach.

the community. The community will hence tend to adopt such a concept. (Sosa 1988, pp. 152–153)

But there are, we think, a number of important unexamined issues that arise from the above proposals—concerning the nature of the notion of normality and the factors that ground it, how precisely normality might fix the threshold for knowing, and finally the consequences that might arise from taking it to fix this threshold. Our aim in the remainder of this paper is to delve more deeply into these issues. Ultimately, we aim to develop the basic idea that knowledge must be sensitive to our epistemic ‘normal’ into an explicit general constraint on knowledge that we call *normalism*.

Normalism is the view that the degree of justification or epistemic support required for knowledge is simply *a normal degree*. As such, it is tailor-made to avoid overdemandingness. If all one needs for knowing that P is *a normal degree* of epistemic support for believing that P, then it is more or less guaranteed that normal people are in a position to know many or even most of the things that concern them in their everyday lives, such as where they live, where to find water and food, or who their parents are. And thus, it is more or less guaranteed that most of their assertions, actions, and so forth will be normatively acceptable—thereby enabling the well-functioning of our various knowledge-linked norms.

As noted in the introduction, we intend normalism as a general framework for a theory of knowledge rather than as a theory proper. As such, we remain neutral on what exactly the epistemic support required for knowledge consists in—it may be the evidence one has, or the reliability of one’s belief-forming processes, or some other epistemic property. All our proposal requires is that epistemic support comes in degrees, which is true on almost any substantive account of epistemic support or justification. In this respect, the normality-constraint resembles other high-level claims about knowledge, such as the contextualist claim that the degree of epistemic support required for knowledge depends on some contextually salient standard (see, e.g., DeRose 1992; Lewis 1996), or the pragmatic encroacher’s claim that the relevant degree of epistemic support depends on what is at stake for the epistemic subject in question (see, e.g., Fantl and McGrath 2002; Stanley 2005). The only additional thing we will assume is that knowledge minimally requires true belief—once again in wide agreement with almost all extant accounts of knowledge.

Normalism specifies a constraint on knowledge which, when incorporated into a full theory of knowledge, would produce something of the following form:

S *knows* that p iff

1. S *believes* that p,
2. p is *true*,
3. S has *at least a normal degree of epistemic support* for believing p,
4. ... (possible additional conditions, such as an anti-Gettier clause).

Conditions 1, 2, and 4 are to be understood in the standard way, with 4 serving as a placeholder for any additional required conditions. The notion of epistemic support invoked in condition 3 may be replaced by whatever account of justification, warrant, or what have you the theorist in question prefers; for example, condition 3 may be used to specify a certain degree of reliability, of evidence, and so forth.

It is the notion of a *normal degree of epistemic support* that constitutes the heart of our normalist proposal; and as noted earlier, it could do with some unpacking. But before elaborating on the particular idea of normality that we will appeal to, it will be helpful to contrast the role normality plays in our account with a few other normality-inspired views within recent epistemology. Within the reliabilist literature, for instance, normality has been invoked to fix the appropriate domain for assessing the reliability of a process. Thus, a process may be held reliable if it does not generate false beliefs under *normal* circumstances (Leplin 2007). A related view appears in Graham (2012), and Smith (2016) uses an explanation-based notion of normality for an account of justification in terms of *normal* evidential support (see also Littlejohn and Dutant 2020). Within the pragmatic encroachment literature, Grimm (2015) suggests that the threshold for knowing will be influenced by the *normal* stakes of persons who might later rely on a given judgment; a similar view appears in Henderson (2009). Peet and Pitcovski (2018), meanwhile, propose that knowledge is *normal* belief, arguing that a requirement of said normalcy is that the belief exhibit characteristic properties of truth and safety, and do so in a characteristic way.

The view we are proposing differs from all of these. Rather than taking normality as a factor in defining justification, reliability, or even knowledge itself, we take normality to merely fix the degree of epistemic support required for knowing. Though Grimm and Henderson similarly tie normality to the threshold for knowing, they do so indirectly, via appeal to the normal level of practical interest. The view we aim to explore, by contrast, is not committed to any kind of pragmatic encroachment—nor is it committed to evidentialism, reliabilism, a safety principle, or any other particular epistemic allegiance. It merely claims as follows: whatever form of epistemic support is required to turn one's belief into knowledge, one needs at the very least a *normal amount* of that support.⁴ The view we propose is, essentially, a purely quantitative constraint on knowing. Nonetheless, although our view differs from the accounts just surveyed, we do suspect that some of the consequences of linking knowledge to normality which we will explore later in the paper may also apply to such accounts.

Let's turn now to characterizing normality itself. Our normalist proposal is inspired by the *folk notion of normality*, which has recently also been studied by experimental philosophers (Bear and Knobe 2017). One might initially suppose that the folk notion of normality simply expresses judgments about what is typical or statistically average in a given domain (see, e.g., Ströbner 2015). For example, one might think that people, when asked about the normal amount of TV to watch per day, will just give their estimate of the average amount of TV that people watch per day. But in fact, people's judgments about the average amount and the normal amount of watching TV per day differ substantially—roughly by an hour. Something similar holds for a number of further activities and events, such as the normal amount of exercising per week, or the normal amount of computer crashes per month (Bear and Knobe 2017, p. 28). In addition, Bear and Knobe also asked people about the ideal amount of watching TV per day, of exercising per week, etc., and they found that normality-judgments were

⁴ Though Grimm and Henderson use normality to 'indirectly' fix the threshold for knowing (via stakes), their views will plausibly render similar thresholds to ours in practice, at least on a very large proportion of cases. We do think that there are many scenarios where the two types of view come apart, but an examination of such cases is best deferred until after our proposal has been presented in full.

predicted by *both* people's *descriptive considerations* about the relevant average, and by their *normative considerations* about the relevant ideal. Thus, folk normality seems to depend both on statistical and normative considerations about the target domain (see also Alexander 1973). Since Bear and Knobe also found that there is no straightforward way to “compute” normality values from people's judgments about what is average and what is ideal, they reasonably conclude that the folk notion of normality is “a hybrid of the two” (Bear and Knobe 2017, p. 26), and thus possibly determined by further factors beyond a mere ‘aggregation’ of the average and the ideal (to be investigated in further experimental work).⁵

Our everyday notion of normality therefore expresses a normative standard that is sensitive to, and shaped by, our actual descriptive reality, while not merely echoing it—and *mutatis mutandis* for our ordinary judgments about the ideal in a given domain. In other words, folk normality can be seen as a “*grounded ideal*”—one that neither trivializes nor vitiates our relevant aims and achievements under ordinary, realistic conditions. For these reasons, we claim that folk normality gives fans of the pedestrian approach exactly the right sort of normative standard for knowledge—one that pushes us to do better while nonetheless respecting our limitations. Moreover, the proposed constraint on knowledge in terms of folk normality could be easily applied by the folk themselves, which also makes it plausible that folk ascriptions of knowledge are *already* sensitive to considerations about normal degrees of epistemic support—at least to the extent that the folk are not influenced by potential infallibilist strands in our folk concept of knowledge (cf. Nichols and Pinillos 2018). In other words, we find it plausible that folk normality provides an important element of a descriptively adequate account of folk attributions of knowledge. We also hold that it provides a compelling starting point for a prescriptive account of knowledge—though a prescriptivist might want to argue for some amount of deviation from the particular weighting the folk give to the average versus the ideal.⁶

We should note that there may be *other senses of normality* in our folk usage of ‘normal’ that work differently from the above account (inspired by the empirical work of Bear and Knobe 2017). For example, in the philosophical literature one finds senses of normality that are frequentist or statistical (see, e.g., Ströbner 2015), or grounded in typical or characteristic properties (see, e.g., Peet and Pitcovski 2018), or species- or kind-relative (see, e.g., Schurz 2001). We think these are all perfectly fine notions of normality for specific (theoretical) purposes, and they may indeed play a role in ordinary folk usage of ‘normal’ in various contexts. However, we also think that the notion studied by Bear and Knobe (2017) captures a core use of the folk notion of normality, namely, the notion of normality *simpliciter*—in contrast to various other uses of normality *for a given thing or purpose* (see also below). Normality judgments based on characteristic properties (such as the property of laying eggs for a chicken, or of having stripes for a tiger), for instance, inherently involve the explicit specification of a reference class (laying eggs is normal *for chickens*, e.g.). By contrast, we claim that the Bear and Knobe notion is something like the *default notion* of folk normality, which

⁵ A recent study by Wysocki (ms) basically confirms the findings of Bear and Knobe, but also adds some complexities about how statistical and evaluative considerations influence folk ascriptions of normality.

⁶ We won't attempt here to discuss the specifics of said weighting; more empirical work will be needed before the folk notion of normality is fully understood.

the folk routinely employ in the absence of further qualifications. Moreover, Bear and Knobe review evidence indicating that even judgments about what is characteristic or prototypical of a given category are influenced by both statistical and normative information. This suggests that there really is something *fundamental* about the notion of normality that we work with in this paper.

Now that the basic notion of normality that we will be using is in place, a few preliminary points should be made about our specific formulation of the normality condition on knowledge. First, note that the relevant degree must be *at least* a normal degree of epistemic support; this is meant to acknowledge that higher degrees of epistemic support beyond a normal degree are, of course, sufficient for knowledge as well (provided that the other necessary conditions are in place).

Second, we claim that knowledge requires at least *a* normal degree of support, rather than at least *the* normal degree of support. The reason for this is that there is no single 'normal'—normal values tend to occupy ranges that lack precise lower (or upper) bounds. That is to say, like many folk concepts, 'normal' tends to incorporate a certain amount of vagueness and slack.⁷ For example, it would be quite strange to claim that the normal amount of exercise per week is exactly two hours. More realistically, the normal amount of exercise per week would fall in a (somewhat vague) range of, say, two to three hours per week.

One might worry, however, that the vagueness and slack in our folk notion of normality is more severe than this example indicates—and to a degree that would be troublesome for our normalist proposal. For example, one might worry that our folk notion of normality will be *wildly context-sensitive* in a way that makes its ordinary usage too unstable for epistemological purposes, and one might also worry that there will be *too much individual variation* in judgments and standards of normality.

We will address the question of context-sensitivity shortly. Concerning the individual variation worry, this is simply not supported by the available empirical evidence: people's normality-judgments exhibit a pattern that shows a remarkable stability and coherence across different domains and several studies with different methodological approaches (cf. Bear and Knobe 2017). We conclude that, at present, there is no good empirical reason to believe that our folk notion of normality might be problematically subject to massive individual variation. At most, folk-normality likely has just about as much vagueness and slack as other notions that are routinely employed in epistemological theorizing, such as justification or reliability.

Third, on the notion of normality just presented, it is plausible that normal values will always be greater than zero—thus guaranteeing that mere true belief never suffices for knowledge. This is because normality is a function of *both* the statistical average and the ideal. Even if the average degree of epistemic support for believing P were to be zero, the ideal presumably will never be. Therefore, even in such zero-average circumstances the normal will be pulled towards the ideal, and thus at least some distance away from zero. Indeed, as we'll soon argue, we find it plausible that the

⁷ We leave it open that a prescriptive account of knowledge might specify a precise lower bound, eliminating this vagueness.

ideal will always be quite high. Thus, we can arguably expect the lowest bound for normality to always involve some non-trivial amount of epistemic support.⁸

Let's turn now to the factors that determine what is 'normal'. Normality, on the view we are presenting, is a function of statistical averages and normative ideals. But average for whom? Ideal for whom? Even when no reference class is explicitly specified—that is, when we consider what is normal *simpliciter*, rather than normal for such-and-so person or group—*some* reference class must be utilized in order to fix the relevant average/ideal values. One possibility, following contextualist views, would be to claim that the standard for 'normal' (and thus the threshold for knowing) is set by the *circumstances of the ascriber* of knowledge. Thus, the relevant statistical average would be calculated over the epistemic states possessed by the ascriber (perhaps over some specified time period such as their adult lifetime); the relevant ideal would be indexed to the ascriber as well. An alternate possibility, following many pragmatic encroachment views, would be to claim that the standard is taken from the *circumstances of the (epistemic) subject*. In fact, we think that neither of these will work; though normalism, like these views, allows the threshold for knowing to be sensitive to contingent facts about the relevant circumstances, the sort of 'indexing' employed by normalism will need to be quite different.

The trouble is as follows. Interpreting a 'normal' degree of epistemic support as 'normal for the attributor' or 'normal for the subject' threatens to make knowledge ascriptions implausibly relativistic, since what is normal for a person A and what is normal for a person B could be very different indeed. An attributor who applies her own 'normal' when ascribing knowledge risks being overly egocentric—just as a TV critic would be overly egocentric were she to use the normal amount of TV-watching *for her* as a standard by which to judge others' normalcy. And using a standard calibrated to any single individual, either ascriber or subject, risks being overly permissive—a very epistemically unsuccessful subject might have an extremely low average degree of epistemic support, such that even after the influence of the ideal the threshold of '*normal-for-them*' would be intolerably low. Ultimately, although some contextual variation is unproblematic, an assertion practice, e.g., that is regulated by droves of widely varying, individualized standards is not likely to be well-functioning. What is needed is a standard that is calibrated to some larger group. And indeed, this is what we find in folk-judgments about what is 'normal' *simpliciter* (as opposed to judgments about what is 'normal *for*'). For example, when people answer the question what a normal amount of watching TV per day is—without 'normal' being in any way qualified or relativized—they seem to be latching onto the prevalent standards of normality in their larger society or culture, instead of some more localized or relativized standard (cf. Bear and Knobe 2017).

As noted earlier, authors like Sosa have invoked the notion of 'normal for a community'. This approach has the virtue of relativizing 'normal' to a group rather than to an individual, thus avoiding the egocentrism worry above. However, there is much room for debate over where exactly the boundaries of the relevant community should be fixed. It is plausible, we think, that standard folk ascriptions of knowledge appeal

⁸ And, of course, on a prescriptive approach to knowledge, one may simply specify that the ideal is to be weighted sufficiently to rule out a zero-valued normal.

to a group to which the ascriber belongs—their community, which will typically, but not always, include the subject. Just as with morality, we don't often judge others by the standards of an out-group. We also think it is plausible that the boundaries of the relevant community are not, say, national or cultural; they are likely far broader. We speculate that they are set by shared epistemic features—in other words, that the relevant community for knowledge ascriptions is an 'epistemic community'.

We won't here fully define what constitutes an 'epistemic community'—doing so would require much more space than we have here. But we will take a stab at narrowing down the options. What we have in mind is a group that has *roughly similar epistemic resources*, such as available cognitive processes, access to information, and so forth, and whose members have *mutual epistemic access* to one another (via verbal communication, written texts, and so forth). For instance, an earthworm is not a member of our epistemic community due to its impoverished cognitive abilities. Medieval serfs are also out, due to their vastly inferior access to education and information, and the lack of mutual epistemic access. And aliens in distant galaxies are out simply because we lack epistemic access to them. A comfortable majority of present-day humans, however, are plausibly within our epistemic community, despite moderate differences in, e.g., intelligence or education.

The driving thought here is that if the reference class determining 'normal' is too narrow, this would potentially interfere with the well-functioning of evaluative practices invoking knowledge as a norm. If I were to take my epistemic community to consist only of highly intelligent PhDs with luxurious amounts of free time for thinking, for instance, the resulting epistemic standard would be unfair when applied during interactions with my overworked mailman or grocery store cashier. Similar issues could arise for a too-broad reference class (though they seem comparatively less likely).⁹ The notion of an epistemic community as roughly defined above strikes us as pinpointing approximately the correct scope. Beyond these rough remarks, insofar as our aim is to describe how people *actually* ascribe knowledge, the exact nature of an epistemic community is simply an open empirical question. One could also approach this issue prescriptively, by asking what the most appropriate boundaries for an epistemic community *would* be; this is a fascinating question, but one which we unfortunately don't have space to pursue here. In any case, the notion of an epistemic community as outlined above provides a relatively 'broad' (but by no means universal!) reference class for determining what counts as a normal degree of epistemic support.

Normalism thus builds in some amount of sensitivity to the attributor, by indexing normality to the attributor's epistemic community. We should emphasize, however, that this is not the sort of sensitivity that is suggested by bank cases and their ilk: it is not, e.g., a sensitivity to stakes or practical interests. Instead, it is a sensitivity to the contingent, *a posteriori* facts about the average epistemic standing of members of the relevant community. 14th century Europeans in the midst of the Black Death faced extremely high stakes with regard to propositions concerning, for instance, effective protection from contracting the bubonic plague; nonetheless, on our normalist account, the standard for knowing said propositions (in that community) was plausibly

⁹ A potential example: suppose hyper-intelligent alien species were discovered to exist. Including their epistemic states in the reference class that determines our 'normal' might shift the threshold for knowing into a range that is overdemanding for humanity.

lower than it is at present. This is due to the fact that lower levels of education and informational access meant that the average level of epistemic support was quite low.¹⁰ Normalism, then, is not best viewed as yet another variant on contextualist or pragmatic encroachment accounts; the ‘normality’ facts that set the threshold for knowing are derived from statistical facts about the epistemic states of the relevant community, along with facts about normative epistemic ideals of the community. Though normalism is thus sensitive to certain practical facts (e.g., average availability of information), these facts are not at the level of the practical context of an individual subject or attributor.¹¹

Nonetheless, we take no stance on the overall influence of individual stakes or practical interests on knowledge-ascriptions. We are open to the idea that they, too, might affect the threshold for knowing, perhaps in individual cases dragging that threshold away from the ‘anchor point’ set by the epistemic community’s normal. Thus, we want to leave open (but not outright commit ourselves to) the possibility that the truth of a knowledge judgment is not just determined by the normal amount of epistemic support for a member of the relevant epistemic community, but rather by the normal amount of epistemic support for a member of the relevant epistemic community *in similar circumstances*, where these circumstances might include practical stakes. As a final speculative possibility, there might be cases in which *both* a broader circumstance-invariant *and* a narrower circumstance-relative reference class combine—or even compete—to determine our knowledge judgments.

For simplicity’s sake, in this paper we’ll content ourselves to defend a more straightforward version of normalism, where normality is determined, for a given ascriber A, by a single and relatively broad reference class consisting of the epistemic community to which A belongs. Our proposal, then, is that a normal degree of epistemic support will be determined by some hybrid of what is average and what is ideal in A’s epistemic community, and thus that the standard for knowing is relative to the epistemic community of the knowledge-ascriber.

Once the relevant community is fixed, the route to determining the average level of epistemic support in that community is at least in principle straightforward—it is

¹⁰ This is one sort of case where our account potentially renders different verdicts than those of Henderson and Grimm mentioned earlier. An account which bases the knowledge-threshold wholly on the normal stakes of a community would suggest that the normal *stakes* for 14th century Europeans regarding plague-related propositions would raise the knowledge-threshold higher than it is for current Westerners (whose normal stakes are much lower, given the scarcity of bubonic plague cases and the availability of effective treatment with antibiotics within our present environment). A similar divergence would arise for other ‘high-stakes-low-average’ cases. Conversely, ‘low-stakes-high-average’ cases (like, for instance, a society of scholarly gentlepersons of leisure) would, if the threshold were fixed primarily or wholly by normal stakes in a community, result in a low threshold for knowing; our view would suggest a high threshold. Ultimately, however, we suspect that Henderson and Grimm would be open to the suggestion that *both* normal stakes *and* normal levels of epistemic standing have effects on how the threshold is set. Indeed, as an anonymous reviewer suggested to us, the two sorts of normality may interact—rising normal stakes will likely lead to rising normal levels of epistemic standing. As we note below, we are neutral on the effect of stakes, and so we are also amenable to a ‘hybrid’ view of some sort.

¹¹ Normalism would fall under what Michael Hannon (2020) has recently labelled ‘pure(ish)’ epistemic views, by contrast with ‘purist’ views which deny a role for the practical in characterizing knowledge, and with standard ‘impurist’ views that both accept a role for the practical *and* hold that role to involve an individual subject’s practical situation. Thanks to an anonymous reviewer for bringing this paper to our attention.

simply an empirical fact. But what of the ideal? One might worry that the determinants of the ideal component in folk normality have not yet been sufficiently articulated by our proposal thus far. However, for descriptive approaches to knowledge, we think that the ideal, too, is simply grounded in yet-to-be-uncovered empirical facts—facts about the attitudes and values of the ‘folk’. Nonetheless, we can at least say a bit more about what the ideal that informs folk normality is an ideal *of*. The value of the ideal seems to depend on the property or thing that normality is being ascribed to. For example, consider the example of a normal amount of exercising per week from Bear and Knobe (2017). In our community, which sets the relevant ideal here, exercising is regarded as a very good or perhaps even essential thing to do for maintaining one’s health and physical fitness. Therefore, it seems plausible that the ideal amount of exercise would remain significantly above zero even if we collectively turned into a society of couch potatoes who don’t exercise at all.

In other words, the ideal relevant to normality is derived from how good or normatively essential we take the target to be, and this can differ quite a bit from one thing to another (e.g., the situation might be very different when it comes to the normal amount of watching TV per day, because watching TV is typically not seen as an especially good or normatively essential thing to do). Applied to our normalist proposal, we claim that positive epistemic support is regarded by our community as a very good or normatively essential thing for forming beliefs about the world. We recognize, for instance, that low levels of epistemic support tend to lead to errors, which in turn lead to negative practical consequences. Our attitudes towards positive epistemic support are rather more like our attitudes towards exercise than our attitudes towards TV-watching. Thus, it seems to be a plausible prediction that the relevant ideal will remain significantly above zero—even were we to turn into the cognitive analog of a society of couch potatoes.

This last observation also helps to answer a related worry, namely, that the relevant ideal might be too liberal in a society that systematically believes things on the basis of very poor epistemic support. We think that this won’t be so, because our ideal of epistemic support as being normatively essential for forming beliefs about the world will keep a normal amount of epistemic support clearly in the positive range, even if our average amount of epistemic support should be very low indeed. Any society that values the practical benefits of true belief will, we think, place the ideal corresponding to knowledge fairly high. Again, however, we stress that this is an empirical matter and the specifics must ultimately be settled by further investigation of the nature of folk normality.¹²

Let’s close this section by looking at a particularly interesting consequence of taking knowledge to be dependent on normality in the way we have specified. On this way of cashing out normality-dependence, the extension of ‘know’ might shift, e.g., over time or across communities. For example, a citizen of ancient Rome might speak truly when he says ‘I know that P’, while a citizen of present Rome might speak falsely when he says ‘I know that P’, only due to differences in the normal degree of epistemic

¹² Thanks to an anonymous reviewer for prompting these clarifications. We have focused in this segment on the ideal corresponding to descriptive approaches to knowledge, which we take to be set by the attitudes of the folk. For prescriptive approaches to knowledge, the appropriate ideal is a matter of philosophical debate—and arguing for any particular standard is unfortunately out of the scope of this paper.

support for believing that P in the ancient Roman's epistemic community and in the present epistemic community of citizens of Rome. Note, though, that the same citizen of present Rome might also speak truly when he claims that the ancient Roman did not know.

The normality-constraint thus leads to a form of *context-dependence* of knowledge-claims; but the relevant 'context' is a lot more stable and less volatile than in standard versions of epistemic contextualism, where the (conversational) context at issue may change quite easily and within seconds, e.g., by walking from the street to the epistemology seminar (see, e.g., DeRose 1992; Lewis 1996). In contrast, for significant changes in what is normal to come about, there would have to be significant changes in what is average or what is ideal in the relevant reference class, i.e., the epistemic community of the ascriber—and such changes usually take time and require large-scale developments, e.g., in the relevant social or natural environment.

This 'slow-motion' context-dependence of normality strikes us as a benefit, because it makes knowledge-ascriptions adaptive in the face of major changes of people's epistemic circumstances, yet without making them unduly volatile. For example, consider societies that move from conditions of informational scarcity to informational abundance, just as has happened in the last few centuries in the Western world with the invention of printing in the 15th century, the rise of electronic mass media in the 20th century, and the expansion of the Internet in the last few decades. For the anti-skeptic in particular, it should be good news that the epistemic support required for knowledge is sensitive to such large-scale epistemic revolutions, because otherwise knowledge might either be too hard to acquire in information-scarce societies to fulfill its everyday role, or it might be too easy to acquire in information-rich societies, such as our own, and thus be trivialized as a normative standard. Instead, what counts as knowledge changes as our epistemic community's circumstances change, and so the extension of 'knowledge' always remains roughly calibrated to said circumstances, and thereby fixes an appropriately demanding normative standard for permissible assertion, action, and so forth.

4 Normalism at work

Let us now consider how normalism—as we have presented it—helps us to both understand and vindicate the many everyday roles that knowledge arguably plays in our lives.

First, let us briefly consider the issue of *skepticism*. On the normalist proposal, knowledge only requires a normal amount of epistemic support. Thus, ordinary people in normal situations will have most of the knowledge they take themselves to have; as advertised, the normalist proposal makes knowledge pedestrian. But despite the relatively easy availability of knowledge on the normalist proposal, normalism still does not make knowledge a trivial epistemic achievement. For one thing, true belief itself is not always easy to achieve. For another, due to the influence of the 'ideal', the normal degree of epistemic support will usually be substantially higher than just any amount of positive epistemic support (as argued in the previous section). Therefore, people can and regularly do fall short of having a normal degree of epistemic support

for their beliefs, in line with the platitude that people—although they know lots of things—also fail to know, or are ignorant of, a fair number of things.

Normalism doesn't rule out skepticism *by definition*, however; a complete lack of knowledge is still possible in very extreme epistemic circumstances or in wholly epistemically intractable domains. For example, imagine a society that falls into a state of complete 'epistemic chaos', where the amount of epistemic support that is available in most domains is very small or even zero (we might imagine, perhaps, a virulent epidemic that causes the destruction of all perceptual capacities along with total amnesia). In such a society the average person might indeed be in a skeptical or near-skeptical epistemic state. For, although the average degree of epistemic support would be very small or zero in such circumstances, the influence of the ideal would nonetheless ensure that the threshold for knowing requires non-trivial epistemic support. Knowledge might, then, be out of reach for members of this epistemically impoverished community, and thus, the vast majority of their assertions and actions might indeed be unwarranted. A functional practice of, e.g., assertion might well be impossible in such utterly dysfunctional circumstances.

A more realistic scenario is 'partial skepticism' for specific domains, for example, concerning propositions about the distant future. It is commonly accepted that we do not know (and also lack non-trivial epistemic support for) propositions about how our society or natural environment will look like 5 million years from now (if it still exists!). Normalism captures this commonplace by holding that our actual degree of epistemic support for propositions about our surroundings in 5 million years is smaller than a normal degree of epistemic support, due to the influence of the relevant ideal degree of epistemic support on normality. In sum, normalism gives us the robust, realistic anti-skepticism that most ordinary people—and most contemporary epistemologists as well—simply take for granted. What is more, normalism also explains *why* robust anti-skepticism is true, and how certain forms of partial, domain-specific skepticism can be true as well.

Normalism has a further advantage for fallibilist, anti-Cartesian accounts of knowledge: it provides an elegant and non-arbitrary solution to the so-called 'threshold-problem' (see, e.g., Hannon 2017, 2019). Unlike infallibilism about knowledge, which sets an extremely high but principled standard for knowledge, issues are less clear for the fallibilist. Suppose you are an evidentialist of some sort about epistemic support. Then, while the infallibilist has a clear and non-arbitrary standard for knowledge, namely, that your evidence needs to exclude all possibilities of error, it is much less clear where the fallibilist should draw the line. In other words, if fallibilist knowledge does not require that your evidence excludes *all* possibilities of error, then *how many* does it have to exclude? The normalist answer to this question would be that knowledge requires *a normal amount of evidence* as epistemic support, such that one's evidence has to exclude *a normal amount or range of error possibilities*. In this way, normalism provides an independently motivated and therefore non-arbitrary solution to the threshold problem. It may still be somewhat vague how much exactly a normal amount of evidence is, but this kind of vagueness is only to be expected for ordinary, non-technical notions like KNOWLEDGE.

As we've noted, normalism's driving motivation is to secure the efficacy of our knowledge-linked practices, such as assertion, action, and so forth. The fact that

normalism ensures that knowledge is pedestrian contributes to this aim. But the context-dependence that we discussed in the previous section has fortuitous implications here, too. Consider, first, how normalism interfaces with the *knowledge-norm of assertion*, i.e., that one should assert that P only if one knows that P. We've already noted that a skeptical epistemic standard, such as certainty, would undermine the practical utility of this norm. Since we are rarely, if ever, certain of the propositions we believe, almost all of our assertions would then be unacceptable. It is also easy to see how normalism avoids this dire outcome, because for humans a normal degree of epistemic support will typically be much lower than certainty or infallibility, and will often be relatively easy to attain.¹³ But normalism's advantages here extend beyond simply implying that knowledge is pedestrian. Normalism's context-dependence also prevents problematic inflexibility, in the following sense: if knowledge were *insensitive* to substantial changes in an epistemic community's situation, the knowledge-norm of assertion could become non-functional for societies that are subject to substantial (but non-total) epistemic decline—perhaps as the result of some apocalyptic scenario, which deprives the society of a large proportion of its informational resources and scientific expertise (say, a widespread nuclear war that destroys internet infrastructure, all libraries, and most of the world's population). Normalism's adaptability leads to the sensible prediction that an appropriately large fraction of assertions will still be normatively appropriate for members of that epistemic community even after the 'informational apocalypse', since the extension of 'know' will change to reflect the new normal.

Note that we should distinguish between this scenario and the near-skeptical scenario discussed earlier; in the latter but not the former, the epistemic loss exceeds the limits of the knowledge-concept's flexibility. The exact point at which this occurs is an empirical one, dependent on the actual degree to which the ideal affects the threshold for knowledge; the cases given, however, strike us as plausible examples of what societies on different sides of the 'line' might look like.

Similar considerations apply to the *knowledge-norm of action and practical deliberation*. Here, too, it is easy to see that, e.g., the epistemic standard of certainty would completely undermine the normative standing of almost all of our actions and practical deliberations, because it is near-impossible for humans to carry out actions or practical deliberations that are only based on propositions believed with certainty. But equally, an epistemic standard that is insensitive to major changes in one's epistemic circumstances, as in the previous example, might undermine the normative standing of everyday actions and practical deliberations. Normalism, by contrast, avoids this unhappy consequence, because the normality-standard adapts to major changes in our epistemic situation.

Let us now consider the *knowledge-norm for everyday inquiry*. It is again easy to see how a skeptical standard would render this norm pointless. But on the common assumption of a fixed, inflexible epistemic standard for knowledge, massive changes in the epistemic environment could also have problematic implications. Consider the example of an everyday inquiry into when one's train departs. One hundred years ago,

¹³ In contrast, certainty might be normal and therefore an appropriate normative standard for angels or god-like creatures. The normalist proposal thus has the plausible implication that god-like creatures should only assert that P if they are certain that P.

it was good enough to carefully read the printed timetable at the train station in order to ascribe knowledge about the departure time to oneself, and thus be normatively entitled to terminate one's relevant inquiry. Today, however, one would arguably have to check the electronic timetable at the train station or real-time information on one's smartphone, even if one has already carefully read the printed timetable; after all, given the ease of checking the more up-to-date electronic information, it would seem 'epistemically lazy' to fail to conduct this simple bit of further inquiry. Only after this 'double-checking' would a 21st-century-traveler be entitled to ascribe knowledge about the departure time to herself, and thus be entitled to terminate her inquiry.¹⁴

Unless the normative standard for knowledge is sensitive to substantial changes in the epistemic environment, it is difficult to accommodate the intuitive normative difference between a traveler in, say, 1919 and 2019. After all, both of them have access to a printed timetable at the train station, yet while solely consulting the printed timetable is a sufficient epistemic basis for the traveler in 1919 to terminate her inquiry into the train's departure time, it is not for the 2019 traveler. The normalist proposal makes straightforward sense of this difference, because the normality-standard adapts to the fact that 2019 travelers have access to more, and more accurate, information about train departure times than 1919 travelers. In normalist lingo, while merely checking a printed timetable at the station is a normal amount of information about train departure times in 1919, it is clearly not a normal amount in 2019.¹⁵

As noted in Sect. 2, it has frequently been argued that *knowledge sets the normative standard for testimony*, or, at least, that *being in a position to know is required for being a good informant*, assuming that good testimony is also possible in the absence of knowledge (cf. Lackey 1999). Similar considerations apply here as in the case of the knowledge-norm of assertion.¹⁶ In the state of an 'epistemic apocalypse', as discussed above, an inflexible standard for knowing might render a knowledge-norm on testimony and good informants pointless, because this norm would fail to distinguish testimony and informants that are acceptable even in such dire circumstances from those that are unacceptable. For example, the testimony that cholera is spread by dirty water seems permissible in the epistemic apocalypse, even if it is only based on a vague feeling that results from subconscious memories. This permissibility is something that a knowledge-norm insensitive to major epistemic changes cannot accommodate. Normalism, on the other hand, can do justice to the normative appropriateness of testimony on the basis of relatively weak evidence in epistemically desperate circumstances.

Once again, similar points hold for *knowledge as an epistemic norm for moral responsibility, praise, and blame*. For example, suppose a traveler S1 in 1919 needs

¹⁴ Assume for the sake of the example that present trains and trains one hundred years ago are equally reliable; in fact, modern trains are likely *more* reliable, but our intuition is that one still ought to 'double-check' with the electronic schedule before ceasing inquiry.

¹⁵ To reinforce this point, consider certain other practices of inquiry that now seem obligatory—for instance, looking information up (possibly using multiple sources) on the internet instead of in a decades-old printed encyclopedia, or getting updated news about a very volatile current event online rather than through a print news source. In general, given the easy availability of up-to-date information in our current epistemic environment, relying solely on print media seems to us to be (in many cases) not merely old-fashioned but plausibly epistemically impermissible.

¹⁶ But see Lackey (2008) for non-paradigm cases of testimony that do not involve assertions, e.g., reading someone's secret diary.

to catch a train in order to reach his destination on time, e.g., because he is a doctor and needs to get to an emergency surgery. If S1 can come to know when his train departs simply by carefully reading the printed timetable at the train station, then an invariant epistemic standard would predict that another traveler, S2, in equivalent moral circumstances in 2019 would also know when his train departs solely on the basis of reading the printed timetable. As a consequence, S2 would *ceteris paribus* act just as morally responsible as S1. In fact, however, we would criticize S2 for having only relied on the printed timetable—S2 arguably displayed morally blameworthy negligence. Normalism makes perfect sense of this situation, because it predicts that S2, in his technologically advanced circumstances, *does not know* when his train departs merely on the basis of reading the printed timetable, for the normal degree of epistemic support for propositions of this kind is much higher in 2019 than in 1919. Given that S2 lacks the relevant knowledge,¹⁷ it is understandable why we would criticize S2 for only consulting the printed timetable.¹⁸

5 Conclusion

We have aimed to develop an account of the threshold for knowing that respects the constraints that arise from various normative roles that knowledge plays in everyday life. Our account is anchored by the claim that the folk notion of normality provides the right sort of mix between realism and idealism to ground such knowledge-linked practices as assertion and action. As noted earlier, we view normalism as a framework under which more specific theories might fall, rather than as a full-fledged, fixed account—much in the spirit of such broad theory-types as contextualism or subject-sensitive invariantism. Though we have presented several speculative proposals, we remain open to a variety of possibilities regarding the particular nature of the concept of normality, the sorts of factors to which it is (or should be) sensitive, the kind of reference class it is or should be calibrated to, and so forth. Our primary goal has been to draw attention to the many interesting questions that arise from taking normality to set the standard for knowing.

As noted in the introduction, we think that normalism provides a plausible framework for a theory of knowledge both *descriptively* and *prescriptively*. Well-functioning societies *need* a set of epistemic norms to regulate permissible assertion, practical reasoning, testimony, inquiry, and so forth. We've argued that those norms must not be overdemanding, else they lose much of their practical utility. They must, that is, respect contingent properties about our actual epistemic capabilities and resources.

¹⁷ More precisely, ascriptions of knowledge about the train departure time to S1 by people in his own epistemic community would be correct, while corresponding ascriptions to S2 by people in his own epistemic community would be incorrect.

¹⁸ Normalism implies that, in our own mouths, 'S1 is deserving of criticism' would be true, while the same utterance would be false for S1's contemporaries. We admit that this may seem a bit counterintuitive—we are not strongly inclined to criticize S1, even from our contemporary perspective. We suspect that, in this particular case, a conflicting intuition arises from the 'ought-implies-can' principle; since there was no way for S1 to check a not-yet-invented electronic timetable or smartphone, S1 could not have improved her epistemic position to a level that is standard in our own epistemic community. We suspect that the ought-implies-can principle may trump other considerations in this case.

Fortunately, it seems that our current practices already have this feature. We do not criticize assertions and actions simply because they fall short of certainty or some other exceedingly high epistemic standard. We employ a standard that makes failure the exception, rather than the norm.

Nonetheless, our normative standards should not be, and in fact are not, simply self-congratulatory—people do fail to meet them, quite regularly. This suggests that our actual, descriptive standards incorporate *some combination* of a) *facts about our actual epistemic performance*, and b) *ideals of epistemic performance*. Empirical work suggests that folk judgments of normality in non-epistemic domains employ just such a combination of the actual and the ideal; we speculate, then, that judgments of normality may already be part of the cognitive processing that generates knowledge-ascriptions. In other words, we suggest that normalism provides a promising starting point for any descriptively adequate account of knowledge.

But in addition, we think that normalism can equally serve as a promising starting point for a *prescriptive* theory of knowledge. As we've noted, well-functioning societies need epistemic norms to regulate assertion and all the rest. However, given the points above, we think that a *prescriptive* account of knowledge—one aiming to generate the concept that best fills the relevant normative roles—will plausibly look quite similar to a *descriptive* account of knowledge. In other words: our current concept of knowledge doesn't likely need much improving in this respect! Nonetheless, we think it's an open question whether, for instance, the folk's current mix of actual-to-ideal reflects the most effective possible ratio. Perhaps we could improve our knowledge-concept by increasing our standards slightly—that is, by weighting the ideal slightly more than we currently do. Or, perhaps we would be better served by multiple epistemic concepts with slightly different weightings—one for assertion, one for action, one to feed into moral judgment, and so forth. The possibilities are vast, and go far beyond the rapidly shrinking space remaining in this paper. Therefore, we'll content ourselves with suggesting that both those looking to analyze and those looking to conceptually engineer our notion of knowledge would do well to look to normality for inspiration.¹⁹

References

- Alexander, P. (1973). Normality. *Philosophy*, 48(184), 137–151. <https://doi.org/10.1017/S0031819100060605>.
- Audi, R. (1997). The place of testimony in the fabric of knowledge and justification. *American Philosophical Quarterly*, 34(4), 405–422.
- Bear, A., & Knobe, J. (2017). Normality: Part descriptive, part prescriptive. *Cognition*, 167, 25–37. <https://doi.org/10.1016/j.cognition.2016.10.024>.
- BonJour, L. (2010). The myth of knowledge. *Philosophical Perspectives*, 24(1), 57–83. <https://doi.org/10.1111/j.1520-8583.2010.00185.x>.

¹⁹ For very helpful comments on previous versions of the paper, we would like to thank two anonymous reviewers of *Synthese* as well as Steffen Koch, Alex Wiegmann, and the participants of the EXTRA.4 research colloquium on “Metaphilosophy and Experimental Philosophy” in January 2020. Joachim Horvath's work on this paper was supported by an Emmy Noether grant of the *German Research Foundation* (DFG), project number 391304769. Jennifer Nado's work on this paper was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, project number 13603718.

- Burge, T. (1993). Content preservation. *The Philosophical Review*, 102, 457–488.
- Davis, W. A. (2007). Knowledge claims and context: Loose use. *Philosophical Studies*, 132(3), 395–438. <https://doi.org/10.1007/s11098-006-9035-2>.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929. <https://doi.org/10.2307/2107917>.
- Descartes, R. (2017). *Meditations on first philosophy: With selections from the objections and replies*. In J. Cottingham (Ed.), 2nd edn. Cambridge: Cambridge University Press.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1), 67–94. <https://doi.org/10.1215/00318108-111-1-67>.
- Fassio, D. (2018). Moderate skeptical invariantism. *Erkenntnis*. <https://doi.org/10.1007/s10670-018-0053-1>.
- Fassio, D., & McKenna, R. (2015). Revisionary epistemology. *Inquiry*, 58(7–8), 755–779. <https://doi.org/10.1080/0020174X.2015.1083468>.
- Graham, P. (2012). Epistemic entitlement. *Noûs*, 46(3), 449–483.
- Grimm, S. R. (2015). Knowledge, practical interests, and rising tides. In J. Greco & D. Henderson (Eds.), *Epistemic evaluation: Purposeful epistemology* (pp. 116–137). Oxford: Oxford University Press.
- Hannon, M. (2017). A solution to knowledge's threshold problem. *Philosophical Studies*, 174(3), 607–629. <https://doi.org/10.1007/s11098-016-0700-9>.
- Hannon, M. (2019). *What's the point of knowledge?: A function-first epistemology*. Oxford: Oxford University Press.
- Hannon, M. (2020). Why purists should be infallibilists. *Philosophical Studies*, 177(3), 689–704. <https://doi.org/10.1007/s11098-018-1200-x>.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Clarendon Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571–590.
- Henderson, D. (2009). Motivated contextualism. *Philosophical Studies*, 142(1), 119–131.
- Horvath, J. (2008). Testimony, transmission, and safety. *Abstracta*, 4(1), 27–43.
- Kappel, K. (2010). On saying that someone knows: Themes from Craig. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology* (pp. 69–88). Oxford: Oxford University Press.
- Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51, 227–232.
- Lackey, J. (1999). Testimonial knowledge and transmission. *The Philosophical Quarterly*, 49(197), 471–490. <https://doi.org/10.1111/1467-9213.00154>.
- Lackey, J. (2008). *Learning from words: Testimony as a source of knowledge*. Oxford: Oxford University Press.
- Leplin, J. (2007). In defense of reliabilism. *Philosophical Studies*, 134(1), 31–42.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Littlejohn, C., & Dutant, J. (2020). Justification, knowledge, and normality. *Philosophical Studies*, 177(6), 1593–1609.
- Nado, J. (2017). Knowledge is not enough. *Australasian Journal of Philosophy*, 95(4), 658–672. <https://doi.org/10.1080/00048402.2016.1273959>.
- Nado, J. (2019). Who wants to know? In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 6, pp. 114–136). Oxford: Oxford University Press.
- Nichols, S., & Pinillos, N. Á. (2018). Skepticism and the acquisition of “knowledge”. *Mind and Language*, 33(4), 397–414. <https://doi.org/10.1111/mila.12179>.
- Peet, A., & Pitcovski, E. (2018). Normal knowledge: Towards an explanation-based theory of knowledge. *The Journal of Philosophy*, 115(3), 141–157.
- Rysiew, P. (2012). Epistemic scorekeeping. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 270–293). Oxford: Oxford University Press.
- Schurz, G. (2001). What is “normal”? An evolution-theoretic foundation for normic laws and their relation to statistical normality. *Philosophy of Science*, 68(4), 476–497. <https://doi.org/10.1086/392938>.
- Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford: Oxford University Press.
- Sosa, E. (1988). Knowledge in context, skepticism in doubt: The virtue of our faculties. *Philosophical Perspectives*, 2, 139–155. <https://doi.org/10.2307/2214072>.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Clarendon Press.
- Ströbner, C. (2015). Normality and majority: Towards a statistical understanding of normality statements. *Erkenntnis*, 80(4), 793–809. <https://doi.org/10.1007/s10670-014-9674-1>.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Clarendon Press.

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

Williamson, T. (forthcoming). Justifications, excuses, and sceptical scenarios. In J. Dutant & F. Dorsch (Eds.), *The new evil demon*. Oxford: Oxford University Press.

Wysocki, T. (ms). Normality: A two-faced concept.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.