

A solution to knowledge's threshold problem

Michael Hannon¹

Published online: 7 June 2016
© Springer Science+Business Media Dordrecht 2016

Abstract This paper is about the ‘threshold problem’ for knowledge, namely, how do we determine what fixes the level of justification required for knowledge in a non-arbitrary way? One popular strategy for solving this problem is impurism, which is the view that the required level of justification is partly fixed by one’s practical reasoning situation. However, this strategy has been the target of several recent objections. My goal is to propose a new version of impurism that solves the threshold problem without succumbing to these criticisms.

Keywords Knowledge · Impurism · Justification · Fallibilism · Threshold problem

1 The threshold problem

How much justification is required for knowledge? If you are a fallibilist, then you think a person *S* can know some proposition *p* even though *S*’s justification for *p* is less than fully conclusive.¹ Although fallibilism is widely accepted, it is surprisingly difficult to describe this non-conclusive level of justification in a clear and non-arbitrary way. Perhaps for this reason epistemologists have been largely silent about

¹ Standard formulations of fallibilism include: (a) one can know that *p* even though one’s justification for *p* is less than conclusive (BonJour 2010: 57); (b) the level of justification requisite for knowing that *p* is compatible with *p* being false (Reed 2002: 144); and (c) one can know that *p* on the basis of evidence that does not guarantee the truth of one’s belief that *p* (Brown 2014: 179). Although it is controversial how to formulate fallibilism precisely (see Reed 2002), these standard formulations all give rise to the threshold problem on which I focus.

✉ Michael Hannon
mhannon09@gmail.com

¹ Department of Philosophy, Fordham University, Collins Hall, 441 E. Fordham Road, Bronx, NY 10458, USA

how strong the justificatory component of knowledge must be.² As Laurence Bonjour writes, “it is fair to say that nothing like a precise specification of this [level of justification] has ever been seriously suggested, let alone more widely endorsed” (2010: 61). Stephen Hetherington says this is “a serious problem about the nature of knowledge, one upon which there is scant epistemological comment” (2006: 41).

These concerns give rise to *the threshold problem* for fallibilist accounts of knowledge. This is the problem of how to provide a plausible account of what fixes the threshold (level, degree) of justification (evidence, probability, warrant, supporting ground) for knowledge.³

Expanding on his earlier concern, Bonjour says it is unclear what sort of basis or rationale there might be for fixing this level of justification in a non-arbitrary way. Hetherington agrees: “there is no non-arbitrary answer to the question of how much justification is the minimum amount needed for knowledge” (2001: 144). It is also unclear why *any* level of justification that is less than fully conclusive would have the significance that makes knowledge valuable. Increasingly high levels of justification are valuable because they improve our cognitive situation by making it more likely that our belief is true, but the idea that there is some specific (or “magic”) level of justification that transforms our cognitive situation from not-knowing to knowing seems peculiar. Why would *this* level of justification make such an important difference?

We can put this point differently by thinking of probability as measured by the use of numbers in the interval $[0, 1]$ on the number line. A probability of 0 means that the claim is guaranteed to be false and a probability of 1 means that the claim is guaranteed to be true. Framed this way, we can ask: how probable must your belief be to qualify as knowledge? There are two obstacles to answering this question: first, any point lower than 1 seems arbitrary (why pick that point precisely?); second, it is unclear why achieving some specific level less than 1 would make an important difference in our cognitive situation (by taking us from not-knowing to knowing).⁴

Unless fallibilists are able to answer the threshold problem, the plausibility of their view is seriously in question. According to Hetherington, the standard

² This debate presupposes that the concept of justification is conceptually prior to the concept of knowledge and can be understood independently. For a contrary view, see Williamson (2000).

³ To keep things simple, I’ll speak in terms of the level of *justification* required for knowledge, rather than in terms of evidence, probability, warrant, or supporting grounds. I do not deny that these terms can fruitfully be given distinct senses, but these distinctions do not matter for my purposes. I use ‘justification’ generically, encompassing both internalist and externalist conceptions.

⁴ Is the threshold problem merely rehearsing the familiar problem of vagueness? Both problems have a similar structure. For example, it is plausible that there is no precise point at which a person goes from being not-bald to bald (the property ‘bald’ is vague); similarly, it is reasonable to suppose that a small increase in probability (e.g. 0.01) can never make the difference between not-knowing and knowing (the property ‘know’ is vague). Thus, an inability to locate a precise boundary between not-knowing and knowing would not be a sufficient reason to doubt that some boundary exists. However, it is surely not good enough for fallibilists to say the level of justification required for knowledge is “high” or “strong” or makes one “highly likely” to be right. While an adequate answer to the threshold problem shouldn’t demand more precision than is achievable (or necessary), fallibilists must provide a reasonable degree of approximation. They have not met this challenge, and so I take the threshold problem seriously.

epistemological insistence on their being such a cut-off point “looks optimistic at best, empty at worse” (2001: 146). To avoid this conceptual problem, Hetherington concludes that *any* true belief is knowledge (2001: 145).⁵ Bonjour also says the threshold problem has no plausible solution, but he draws a different conclusion: knowledge is “a myth” (2010: 57). Both conclusions are counterintuitive and should be avoided if possible.

In a recent paper, Brown (2014) considers but rejects one broad strategy for answering the threshold problem, namely, the impurist strategy according to which the required level of justification is partly fixed by one's practical reasoning situation. Brown raises a number of objections against this view, and on the basis of these objections she concludes that impurism does not provide a satisfactory response to the threshold problem.

I want to defend a version of impurism that is importantly different from the type that Brown considers. I will start by summarizing Brown's argument in order to highlight where, on her view, standard formulations of impurism go wrong. I have considerable sympathy with Brown's claims, so I will not dispute them. Reviewing her argument will, however, help to clarify the advantages of my own view. My goal is to show that there is a plausible version of impurism that solves the threshold problem and is not vulnerable to previous criticisms. I call this view *communal impurism*—but I'll say more about this later.

2 Impurism and the threshold problem

One strategy for answering the threshold problem is impurism.⁶ According to impurism, whether a subject S's true belief that *p* is knowledge depends not only on truth-conducive factors but also on the stakes, or how important it is to S that *p* be true (Fantl and McGrath 2002, 2009; Hawthorne 2004; Stanley 2005).⁷ To solve the

⁵ More specifically, Hetherington says mere true belief is a conceptual end point within the graded category that is knowledge: there are a whole lot of potential grades of knowing that *p*, so there is no threshold problem on his view.

⁶ Hetherington (2006) tries to solve the threshold problem in a different way, but there are at least three problems with his solution. First, on his view we can know where the threshold for knowledge is set, but only vaguely and on the basis of very weak justification. In particular, he says that *even if* we know full well (perfectly or conclusively) that any evidence which probilifies my belief to degree 0.9 would be good enough for knowledge, and *even if* we know full well that a probilification of merely 0.3 is not good enough for knowledge, then all we can know is that the justificatory boundary for knowledge falls within that 0.6 span of justificatory support. This is pretty vague and surely not precise enough to meet the challenge of providing a reasonable degree of precision (see fn. 4). Second, if we subtract 0.6 from the maximum 1, we would know that knowledge has a justificatory boundary with only 0.4 strength. But can knowledge of a fact be so poor—so weak? See Hetherington (2001) for a full defense of this idea. Third, Hetherington acknowledges that his solution goes against his own theory of knowledge, as defended in his 2001 book (Hetherington 2006: 45).

⁷ ‘Impurism’ may be understood more broadly as the view that whether some true belief is knowledge does not only depend on truth-conducive factors. My focus is on the more specific view that concentrates on stakes, or how important it is to S that *p* be true.

threshold problem, the impurist argues that the required degree of justification is partly fixed by one's practical reasoning situation.

Hawthorne and Stanley (2008), Fantl and McGrath (2009), and Grimm (2011) implement this strategy. As Fantl and McGrath write,

We hope it is clear how a pragmatist account can help here [with the threshold problem]. How probable must p be for p to be known? It must be probable enough to be properly put to work as a basis for belief and action. (2009: 26)

The impurist solution exploits the sufficiency direction of the knowledge norm of practical reasoning, which states that the minimum level of justification (or probability) required for knowledge is that which is enough for the subject to properly rely on the proposition in her practical reasoning.⁸

Brown distinguishes two ways to implement the impurist strategy, neither of which she finds plausible. They are:

THE UNITY APPROACH: the threshold for a subject to know *any proposition whatsoever* at a given time is determined by a privileged practical reasoning situation she then faces, most plausibly the highest stakes she is then in.

THE RELEVANCE APPROACH: the threshold for a subject to know *a proposition* at a given time is determined by the practical reasoning situation she is then into which that particular proposition is relevant.

Brown attributes the unity approach to Fantl and McGrath (2009) and the relevance approach to Hawthorne and Stanley (2008).⁹ I'll now say a bit more about each view, as well as articulate Brown's objections to them.

2.1 The unity approach

According to the unity approach, there is a single privileged practical reasoning situation facing a subject at a time which sets a general standard for that subject to know *any proposition whatsoever* at that time. This general standard is most plausibly set by the highest stakes practical reasoning situation faced by the subject (I'll explain why in the next paragraph). For example, imagine that Keith and Rachel stop at the bank to deposit their paychecks. They have an impending bill that must be paid and little money currently in their account, so it is very important that

⁸ Fantl and McGrath (2009) and Stanley and Hawthorne (2008) have defended the sufficiency direction of the knowledge norm of practical reasoning; Brown (2008) and Gerken (2015) have criticized it. I will not enter that debate here.

⁹ Fantl and McGrath (2009: 66) write, "If you know that p , then p is warranted enough to justify you in Φ ing, for any Φ ." This formulation imposes no restriction that the proposition p should be relevant to the action in question. Thus, for Fantl and McGrath, the strength of epistemic position required for a subject to know any proposition at a given time is fixed by a privileged practical reasoning situation facing a subject at that time (i.e. the unity approach). In contrast, Hawthorne and Stanley (2008: 578) write, "Where one's choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting if and only if you know that p ." This formulation allows that the strength of epistemic position required for a subject to know one proposition at a given time may differ from the strength of epistemic position required for her to know a distinct proposition at that time (i.e. the relevance approach).

they deposit their paychecks by Saturday. Keith remarks that 2 weeks ago he was at the bank on Saturday and it was open, but Rachel points out that banks do change their hours (DeRose 1992: 913). If the highest practical reasoning situation that Keith faces at 4:30 p.m. on Friday is the decision whether to wait in line at the bank or return tomorrow, then the stakes in this context set the epistemic standard required for Keith to know any proposition at that time. For instance, the same (strong) epistemic position would be required for Keith to know that there are daffodils in the local park (Brown 2014: 181).

Why would the highest stakes practical reasoning situation for a subject set the general threshold for knowledge? The rationale is straightforward: if the privileged practical reasoning situation were identified with a lower stakes scenario, this would undermine the sufficiency direction of the knowledge norm of practical reasoning. It would no longer be true that if a subject knows a proposition, then she is in a good enough epistemic position to rely on it in her practical reasoning. This is a problem for impurism because the sufficiency direction is a premise in one of the most important arguments for impurism (see Fantl and McGrath 2009: ch. 3). If the threshold for knowledge is not set by the highest stakes practical reasoning situation, then Keith could know that the bank is open on Saturday (because he satisfies the threshold set by a lower stakes situation) even though he could not rely on that proposition in his practical reasoning about whether to come back on Saturday.¹⁰ Thus, for the unity approach the general standard for knowledge must be set by the highest stakes practical reasoning situation faced by the subject.

Brown's central objection to the unity approach is that it has unwelcome skeptical consequences. A subject in a high-stakes situation at time t is at risk of knowing very few propositions at t because the epistemic standard set by the subject's high-stakes context at t would also set the threshold for her to know any proposition whatsoever at t . As a result, anyone who is in a high-stakes situation will not only find it more difficult to know propositions related to their practical reasoning environment, *but also to know any proposition whatsoever at that time*. I will call this *The Problem of Skeptical Consequences*.

The impurist might try to confine this skeptical problem to a few points in time; however, Brown argues that this strategy is unavailable because (a) "we are frequently in high-stakes practical reasoning situations" and (b) making high-stakes decisions "often takes a considerable period of time" (2014: 184). In support of (a), Brown points out that everyday life often confronts us with high-stakes decisions. Her examples include "whether or not to enter or end a marriage, to have children, to take a certain job, to move house, to invest one's life savings in certain stocks, and so on" (2014: 184), all of which are plausible scenarios where the threshold for relying on a proposition in one's practical reasoning is quite high. In support of (b), Brown provides the example of deciding whether or not to end a marriage, which may take months or even years. Brown acknowledges that such a person would not be consciously considering this decision at every second of every day over several

¹⁰ As Brown (2014: 183) points out, setting the general threshold for knowledge lower than the highest stakes situation would also undermine Fantl and McGrath's (2009: 22) solution to the problem of concessive knowledge attributions.

months or years, but she claims that “a lack of conscious reflection does not show that one is not in a decision situation” (2014: 184). Brown doesn’t elaborate much on this point, but I will grant it for discussion.

As a result, the unity approach has the consequence that much of what we take ourselves to know is often not knowledge at all. Brown rightly takes this to be problematically skeptical. She therefore recommends that impurists reject the unity approach.¹¹

2.2 The relevance approach

To avoid the skeptical implications of the unity approach, one might turn to the relevance approach. According to this view, the strength of epistemic position required for a subject to know a proposition at a time depends on what practical reasoning situation she then faces to which that proposition is relevant. For example, the fact that Keith is in a high-stakes situation with respect to the proposition that the bank is open on Saturday means that he must be in a very strong epistemic position with respect to this proposition in order to know it. However, the fact that Keith is in this high-stakes situation does not affect the strength of epistemic position required for him to know other unrelated propositions. Although Keith must be in a very strong epistemic position to know that the bank is open on Saturday, this does not entail that he must be in an equally strong epistemic position to know that there are daffodils in the local park.

The relevance approach faces two problems. I will call the first *The Problem of Easy Knowledge*, which Brown summarizes as follows:

The Problem of Easy Knowledge

The [problem] arises from propositions which are relevant only to actions with very low stakes. We may worry that the relevance approach would allow that, in such situations, one can know a proposition on the basis of an extremely weak epistemic position. (Brown 2014: 187)

With Brown, I believe it is problematic to allow that a subject can know a proposition on the basis of an extremely weak epistemic position simply because the specific stakes are so low.¹²

Another worry for the relevance approach is what Brown calls *The Problem of Practically Irrelevant Propositions*:

The Problem of Practically Irrelevant Propositions

The second problem arises from the fact that, for a subject S at a time *t*, there may be some propositions that are irrelevant to any of S’s actual practical decisions at *t*. (Brown 2014: 188)

¹¹ Brown isn’t claiming that the unity approach entails that we know nothing at all. A subject might happen to face no high-stakes situation at a time.

¹² See Russell and Doris (2008) for a similar objection.

To illustrate, Brown asks us to consider the proposition that Mars has two moons. It is true that for me right now this proposition is irrelevant to any practical decisions I face. Granted, it might be relevant at some time to someone else (or to me at some future time); however, the fact that the proposition concerning Mars is practically relevant for, say, a NASA scientist next year is of no help in determining the standards for me now to know that proposition since, for impurists, it is the practical reasoning situation faced by a subject at a time which determines whether she knows.

The relevance approach thus fails to fix the level of justification required to know those propositions that are irrelevant to any practical reasoning situation a subject then faces. This is a problem because we have a vast range of beliefs that are, at any given time, irrelevant to any practical reasoning situation we face. As a result, the relevance approach not merely fails to give guidance for judgment about such cases; it leaves knowledge indeterminate (Brown 2014: 189). This is a significant theoretical cost.

So far my discussion has focused mainly on Brown's argument because it nicely sets up the threshold problem and illustrates why, on the basis of her objections, two current forms of impurism do not satisfactorily answer the threshold problem. In the next section, I will propose a third way to implement the impurist strategy, which I call *communal impurism*. What distinguishes this version of impurism from both the unity and relevance approaches is that, on my view, the threshold for knowledge is not necessarily set by the *subject's* practical reasoning situation. Rather, the practical reasoning situation faced by other inquirers (even potential inquirers) largely determines the level of justification required to know. As I will argue, this view does not succumb to the criticisms Brown raises against the unity and relevance approaches. This is largely because the impurists that Brown has in mind are *subject-centered* impurists, whereas my view is not one that she considers. If we adopt communal impurism, however, we can set the level of justification required for knowledge in a non-arbitrary way, thereby solving the threshold problem.

3 Communal impurism

To support my version of impurism, I will appeal to an increasingly popular idea about why humans think and talk about knowledge. In particular, my strategy is to use Craig's (1990) hypothesis that the function of the concept of knowledge is to identify good informants.¹³ However, my account differs from Craig's in some crucial ways, which I'll highlight as we proceed.¹⁴

¹³ I will speak in terms of *the* purpose (role, function) of the concept of knowledge for expository convenience, but I do not assume that there is only one role played this concept. Rather, I claim this is a central (common, important) role. This hypothesis has been endorsed by Dogramaci (2015), Fricker (2008), Greco (2008), Grimm (2015), Hannon (2014), Henderson (2009), McKenna (2013), Neta (2006), Pritchard (2012), and others.

¹⁴ One important way in which our views differ is that I do not posit a fictional 'state of nature', which is perhaps the weakest and most problematic aspect of Craig's genealogical account. Indeed, I am not providing a genealogy at all. Rather, I claim that the concept of knowledge is the outcome of very general facts about the human situation, especially our need for information from reliable sources to guide our actions.

By Craig's account, the concept of knowledge originates from our need to recommend good sources of information to members of our community. Each of us needs reliable information to successfully guide our own actions, but we also have a common need to pool and transfer information in order to make it easily accessible. It is for our collective benefit that we assess the reliability of informants not just for ourselves but for others, too, since this allows us to store reliable information while it is available, without knowing when, why, or under what circumstances it might be needed. To achieve this goal, we must identify informants who are reliable enough for a variety of individuals with a wide range of interests. To ensure that informants meet this requirement, a practice develops whereby the epistemic standard is set fairly high. According to Craig, this will

edge us towards the idea of someone who is a good informant as to whether p whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability, someone who is very likely to be right—for he must be acceptable even to a very demanding inquirer. (1990: 91)

In other words, a knower is someone epistemically positioned such that many people can freely draw on his or her information and then use that information to guide action.

This hypothesis is based on several plausible assumptions. First, we need true beliefs about our environment in order to successfully guide our actions. Second, we require sources of information that will lead to true beliefs. Third, often the easiest and most efficient way to acquire a true belief is to ask someone reliable. Fourth, on almost any issue some informants will be more likely than others to provide a true belief. Thus, any community may be presumed to have an interest in evaluating sources of information.

The theory of relevant alternatives provides a clear way to articulate this view. Craig does not describe his view in this way, so what follows is my own elaboration on Craig's core hypothesis. I propose the following:

The Communal Threshold for Knowledge

To know that p , an agent must be in a strong enough epistemic position with respect to p to eliminate all of the not- p possibilities that are relevant alternatives to members of the epistemic community that might draw on the agent's information.

This proposal raises a number of questions. First, who are the members of the "epistemic community"? Does it include all possible inquirers? All past, present, and future inquirers? All living inquirers? Second, which alternatives are "relevant" to members of the epistemic community? Third, how should we think of the "might" in the claim that a knower must eliminate the relevant alternatives to members of the epistemic community that might draw on the agent's information?

One thing is clear: to include an unrestricted class of people in the epistemic community would have disastrous results. Skepticism looms if we think that a knower must be in a strong enough epistemic position to satisfy the demands of any inquirer whatsoever, including very demanding inquirers. Sometimes the stakes for an

individual are so high that only the strongest of epistemic positions will meet their demands.

In order to avoid skeptical worries, we should think of the epistemic community as roughly comprised of anyone who might actually draw on the information, where “might” tracks the notion of a possibility that could *reasonably* be expected to occur.¹⁵ This use of “might” does not track mere logical possibility. Suppose I believe that the Red Sox won the baseball game last night because someone reliable told me so. It is widely accepted that reliable testimony typically suffices for knowledge. Now, it is logically possible that someone with life-or-death stakes might need to know whether the Red Sox won last night. Almost certainly, a person with such incredibly high stakes wouldn't regard my second-hand testimony as sufficiently reliable to act on. But the possibility that such a person “might” appeal to my belief that the Red Sox won is not a possibility that could reasonably be expected to occur. The mere fact that such an inquirer could possibly exist does not in general impugn my ability to acquire knowledge through testimony. If fallibilism is true, one need not meet the highest possible demands to qualify as a knower.¹⁶

Another way of putting this, following Grimm (2015), is to say that the threshold for knowledge will reach a level high enough to respect the “typical” or “normal” stakes of people in the epistemic community. Thus, the epistemic community is narrower than all possible inquirers, as well as all actual inquirers. It is comprised of individuals we can reasonably expect to draw on our information.¹⁷ A knower must therefore be reliable enough for anyone who may *reasonably* seek to rely on his or her testimony.

This account helps to clarify which alternatives are “relevant” to members of the epistemic community. The alternatives that must be ruled out are those that are *fitting* or *reasonable* to the members of the epistemic community. Of course, determining the set of relevant alternatives is not an exact science. However, it is clear that we have an idea—rough, to be sure—of what normally counts as having done enough to establish the propriety of a knowledge claim.¹⁸ People develop a conception of how reliable an informant must be to count as a knower because the practice of epistemic evaluation is one that we all grow up into, and which has continued for some time. As a result, we can say tolerably well, in particular cases,

¹⁵ This view is similar to Henderson (2009) and Grimm (2015). Later I will raise some problems for their views and attempt to solve them.

¹⁶ A true belief becomes knowledge when enough justification supports it, but it is still possible that further justification will make one's knowledge better. We needn't insist, as Dretske (1981: 363) does, that somewhere along the spectrum of justification is a point beyond which one's knowledge cannot be further improved by more justification. I merely insist that when a certain amount of justification is acquired, knowledge is achieved and non-knowledge is left behind.

¹⁷ This might plausibly include some non-actual people, as Henderson (2009) argues. Moreover, there are arguably some *actual* skeptical epistemologists who imagine that deceiving demons are relevant to every contingent proposition; thus, it must also be true that the interests of some actual people will not affect the communal standard.

¹⁸ Austin (1946) has a similar view, and Lawlor develops this idea in chapter 5 of her (2014) book. Some experimental epistemologists doubt that there is sufficiently wide cross-cultural agreement to say that “we” have an idea of what normally counts as having done enough to establish the propriety of a knowledge claim (e.g. Weinberg et al. 2001). However, these findings have been disputed by Boyd and Nagel (2014), Hannon (2015), Machery et al. (2015), Nagel et al. (2013), and Turri (2013).

what does and what does not need to be eliminated in order to appropriately credit someone with knowledge. (Here's an example: were I to claim that there is milk in the fridge, I am not open to criticism or blame for failing to rule out the possibility that I am a brain in a vat—at least not in the same way that I could be criticized for failing to take into account the much more everyday possibility that someone finished the milk.) Socialization and acculturation make us proficient at distinguishing the possibilities that must be eliminated in order to have knowledge from those that typically do not. Further, unanimous agreement is neither expected nor required. All we must presume is that such judgments will (or would) coincide sufficiently to give us what Rysiew calls “a set of core not- p alternatives” (2001: 489). Without this assumption, it would be mysterious how people are adept at determining what a speaker means when uttering, “S knows that p .”¹⁹

By reflecting on the epistemic needs of a broad class of people, we find a principled basis for our common judgments about whether a subject knows. As inquirers we want information on which to base our beliefs and actions. Following Henderson (2009), let's call this ‘actionable information’. An agent knows that p if she is epistemically positioned with respect to p so as to be a good source of actionable information, which means that one may take it from her that p . Basic practical concerns thus generate a standard that is fitting to certify actionable information for many people with diverse interests. This sets a general (i.e. communal) threshold for knowledge. The level of justification needed for knowledge is that which puts the agent in a strong enough epistemic position for her to serve as a reliable source of actionable information for many members of her epistemic community.²⁰

However, we also recognize that this general standard is not high enough for every practical reasoning situation. Someone with a particularly important reason for wanting to get the right belief on a question is unlikely to rely on the general public opinion that an individual knows the right answer. DeRose's bank cases nicely

¹⁹ Another way in which my view goes beyond Craig's is that his own description of a reliable informant is not sufficiently detailed to resolve the threshold problem. Craig says that a knower is “someone with a *very high degree of reliability*, someone who is *very likely to be right*—for he must be acceptable even to a very demanding inquirer” (1990: 91, my emphasis). But, as BonJour says, “it is surely not good enough to say merely, as is commonly said, that the level of justification in question is “strong” or “high” or “adequate” or enough to make it “highly likely” that the belief in question is true, for nothing this vague is enough to specify a definite level of justification and a corresponding definite concept of knowledge” (2010: 60). My relevant alternatives account is more precise than Craig's view, but not more precise than the subject matter allows (as Aristotle would say).

²⁰ You might wonder whether the community-wide epistemic standard is the same across all propositions or whether it depends on the proposition (or perhaps the topic) in question. According to the former view, there would be something like an average amount of reliability or justification that is expected for any proposition. I prefer the later view, however, which says the amount of reliability or justification required for knowledge will depend on the relevant proposition (or topic). Although I will not provide a detailed argument for this view here, I will motivate it with an example. In the case of lottery propositions (e.g. ‘I won't win the lottery’), it is implausible to think that we know this even if we have overwhelming statistical evidence to support this judgment. However, if the level of reliability needed to know this proposition were the same for all propositions, then either we would know that we'll lose the lottery or we would know almost nothing (depending on how high you set the general standard). Both results are counterintuitive. Thus, it is more plausible to think that the community-wide standard will depend on the proposition (or topic) in question.

illustrate this point. In the low-stakes context, Keith's claim to know that the bank will be open seems true; meanwhile, it also seems true that in the high-stakes context Keith doesn't know that the bank will be open. This is because in the former context Keith is reliable enough for most people to base their beliefs and to act on his information, but in the latter context his epistemic position is not sufficiently strong to tell against some error possibility that reasonably seems significant in his practical reasoning situation. This illustrates that the individual's stakes may trump the communal standard when it is not sufficiently high for the relevant practical reasoning situation.²¹

Any particular epistemic standard will be too low for some people and too high for others. For this reason, one might worry that my view posits a puzzling asymmetry with respect to whose stakes affect the relevant epistemic standard. On the one hand, I claim that the threshold for knowledge cannot drop below the communal standard because it would cease to serve the function of identifying reliable informants for the community at large. On the other hand, I claim that individual stakes can trump the communal standard when the stakes for an individual are higher than normal (as illustrated in the bank case above). But you might wonder why the individual's interests would ordinarily get swamped by those of the community except when the individual's stakes are exceptionally high.

I can explain this seemingly puzzling asymmetry. On the one hand, the necessity of pooling and sharing actionable information creates a need for a communal threshold for knowledge. This is because sharing useful information requires us to identify informants who are reliable enough for a variety of individuals with diverse interests. To ensure that informants are sufficiently reliable for this purpose, a practice develops of setting the standard fairly high. This communal threshold for knowledge must be high enough to ensure that anyone who meets it will be sufficiently reliable for most practical reasoning situations.²² Moreover, the standard cannot drop below this threshold because to do so would be inconsistent with the function of knowledge ascriptions, namely, to identify people who are epistemically positioned to render information that is fitting to the community's interests. Thus, whatever practical reasoning situations a subject faces, there is a general minimum level of strength of epistemic position required for a subject to know a proposition; this is the communal threshold.²³

²¹ Elsewhere I have called the communal standard the 'default' standard to indicate that it can be overridden when elevated stakes are involved (Hannon 2013).

²² This standard cannot be *too* high because that would make knowledge less than widely available, which would prevent knowledge from playing its role in practical reasoning.

²³ The idea of a communal (or default) threshold might be compatible with insensitive invariantism. According to the insensitive invariantist, what counts as being in a sufficiently good epistemic position to know some proposition does not vary—is not sensitive to—any individual's stakes or practical interests at the time in question, whether it be those of the subject, the attributor, or the evaluator of a knowledge claim. An insensitive invariantist might argue that the communal threshold for knowledge firmly settles at a level high enough to satisfy the function of identifying good informants to the community, and the alleged context-sensitivity of our knowledge ascriptions might be dealt with at the level of pragmatics (see Rysiew 2001; Brown 2006; Gerken 2011). I see no reason to rule out the possibility of such a view. However, my aim in this paper is not to settle the issue of whether we should endorse insensitive invariantism or some rival view. Rather, my aim is to defend a new version of impurism that (a) is motivated by plausible assumptions about the purpose of the concept of knowledge, (b) does not succumb

On the other hand, the communal threshold is not high enough for *all* practical purposes. Sometimes people demand informants who satisfy very high standards, as shown by DeRose's high-stakes bank case. Although Keith meets the communal threshold for knowledge in both the low- and high-stakes cases, his evidence is not strong enough for his practical reasoning situation in the high-stakes scenario. And since an agent knows that p only if his epistemic position is strong enough to tell against some error possibility that reasonably seems significant in his practical reasoning situation, Keith fails to know. When we recognize that serious consequences (and not necessarily our own) turn on having true information, the communal threshold will be inadequate. In this way, my view remains a version of impurism. Keith can know that the bank is open in one context and yet that same proposition fails to be known by Keith in another context, even though the strength of his epistemic position is the same in both cases. Similarly, communal impurism permits a proposition to be known by one person, while that same proposition can fail to be known by another person with the same strength of epistemic position at the same time.

Allow me to summarize this view. Communal impurists claim that the threshold for knowledge is sufficiently high to allow us to pool and share reliable information with a wide range of people that have diverse interests and projects. However, some practical reasoning situations will require us to crank the epistemic standard up at least another notch. Thus, meeting the communal threshold for knowledge is necessary but not always sufficient to know. This account offers a plausible diagnosis of the commonly held intuitions about low-stakes and high-stakes cases. Someone who meets the communal threshold for knowledge will usually qualify as a knower, but she will not qualify as a knower if she isn't reliable enough to meet the more demanding expectations in a high-stakes practical reasoning situation. Unsurprisingly, we should not recommend an informant to an inquirer if we are aware that the inquirer's purposes are particularly pressing and so the informant will fall short of the inquirer's heightened demands. Thus, we do not attribute knowledge in high-stakes scenarios unless a very high epistemic standard is met.

Having outlined communal impurism, let's now re-examine Brown's objections.

4 Revisiting Brown's objections

The Problem of Skeptical Consequences is no problem for communal impurism. We invite unwelcome skeptical consequences if a subject's high-stakes situation at t sets the threshold for her to know any proposition whatsoever at t . But the communal approach is not committed to the idea that a single practical reasoning situation facing a subject at a time sets a general standard for that subject to know any proposition whatsoever at that time. Rather, the threshold for a subject to know a proposition must be at least as high as practical reasoning typically requires, and it

Footnote 23 continued

to Brown's incisive objections against standard formulations of impurism, and (c) offers a plausible solution to the threshold problem.

may be temporarily heightened if an inquirer's practical reasoning situation calls for it. Whether or not the threshold for knowledge is higher than the communal standard will be determined by the practical reasoning situation to which that particular proposition is relevant. Thus, it does not follow that anyone who is in a high-stakes practical reasoning situation will find it difficult to know any proposition whatsoever at that time.²⁴

What about *The Problem of Easy Knowledge* and *The Problem of Practically Irrelevant Propositions*? The first problem arises if a subject can know a proposition on the basis of very weak justification simply because that proposition is relevant to only low-stakes situations the subject faces. The second problem arises if the level of justification required to know a proposition remains unfixed—thereby leaving knowledge indeterminate—because the proposition is irrelevant to any practical reasoning situation the subject faces. Does communal impurism make knowledge too easy or leave knowledge indeterminate?

For communal impurists, the threshold for knowing a proposition does not drop to an implausibly low level (or remain indeterminate) simply because the proposition is relevant only to low-stakes situations (or no practical reasoning situation at all) faced by the subject. Whatever practical reasoning situations a subject faces, the communal impurist claims there is a general minimum level of strength of epistemic position required for a subject to know a proposition. This is the communal threshold for knowledge.²⁵

While I think this account is right as far as it goes, it leaves two issues unresolved.

First, *The Problem of Easy Knowledge* will just resurface at the level of the community. According to communal impurism, the epistemic standard for knowledge is set not just by the subject's practical reasoning situation but also—and perhaps principally—by the practical interests of members of the subject's epistemic community. In other words, the practical reasoning situations faced by other inquirers largely determines the level of justification required to know. However, we can imagine a community in which not much is at stake for *anyone* in the truth of certain propositions. For example, imagine a community whose members are not very interested in baseball results. Communal impurism appears to entail that in this community, one can know that the Red Sox won on the basis of very weak justification. This is a problem because knowledge requires good evidence even with respect to issues in which an entire community might have little stake (i.e. relative to which the practical costs of being mistaken are basically nil). Let's call this *The Problem of Communally Low Stakes*.

Second, *The Problem of Practically Irrelevant Propositions* also resurfaces at the level of the community. According to Brown, impurism fails to fix the threshold for

²⁴ It is theoretically possible for an entire community to be in a high-stakes reasoning situation with respect to some proposition, but that would only make it difficult for the community to know the truth of that proposition.

²⁵ Henderson (2009, 2011) and Grimm (2015) have suggested something similar. As a contextualist, Henderson would articulate his view in terms of 'knowledge', not knowledge. I'll ignore this bit of semantic ascent because I'm interested in knowledge.

knowing propositions that are irrelevant to any practical reasoning situation faced by a subject. Communal impurists like Grimm (2015), Henderson (2009, 2011), and myself (Hannon 2013) have tried to solve this problem by appealing to the relevance of certain propositions to some possible actions that are not currently under consideration. In particular, we have argued that the epistemic standards for knowledge are sensitive to the concerns not just of the subject (or attributor) of knowledge but also to members of the subject's (or attributor's) epistemic community, including merely potential inquirers (who are like us but are interested in the proposition in question). However, some propositions are so trivial that *nobody* in *any* community will likely care about their truth; for example, it is doubtful that anyone does or would care about the number of grains of sand currently on Waikiki Beach. This information will not likely be relevant to any action in which someone has stake.²⁶ Thus, communal impurism seems to provide no guidance in cases where a proposition is of no practical interest to anyone. Consequently, it will be indeterminate what level of justification is required to know these propositions, so it will be indeterminate whether anyone knows such things. This is a problem because even with respect to questions that nobody would reasonably care about, good evidence still seems required for knowledge. Let's call this *The Problem of Communally Irrelevant Propositions*.

When confronted with these two new objections, the communal impurist might simply bite the bullet. She might be willing to do this because her view is already on much better footing than the (subject-sensitive) versions of impurism that Brown considers. How is it on better footing? Well, if communal impurists are right that a knower's epistemic position must be strong enough to meet the community's general minimum threshold, then knowledge would not drop to an implausibly low level simply because some subject does not have much stake in the issue. For instance, although *I* might not care who won the 1990 World Series, it is plausible that *someone* cares about it. Communal impurists thus greatly shrink the amount of easy knowledge in the world, giving their view an advantage.

Communal impurism also shrinks the amount of indeterminacy about knowledge. This is because the threshold for knowing that *p* is not typically set by the relevance of *p* to some subject's practical reasoning situation, but rather by the relevance of *p* to a wide range of possible actions that are not currently under consideration (i.e. to the actions of other people or possibly to oneself at some future time). Thus, communal impurism can provide a fixed threshold for knowing many propositions, including those irrelevant to any practical reasoning situation faced by a given subject. The only propositions that do not have a fixed threshold for knowledge would be ones that we cannot imagine *anyone* reasonably appealing to. In this way, the communal impurist mitigates another of Brown's concerns.

But instead of just blunting Brown's objections, we should try to address them more completely. This requires us to deal with *The Problem of Communally Low Stakes* and *The Problem of Communally Irrelevant Propositions*. In what follows, I

²⁶ I am not suggesting that seeking the truth is always motivated by practical concerns: we may want to know something simply because it satisfies our curiosity. Nevertheless, it is unlikely that anyone would (or should) be interested in the number of grains of sand merely for its own sake.

will sketch an answer to these problems. Getting clear on these issues will not only further distinguish my view from other communal impurists, such as Grimm and Henderson, it will also provide an opportunity to fully address the threshold problem and move the discussion forward.

5 Refinements

Let's start with *The Problem of Communally Low Stakes*. Are communal impurists committed to the view that people in a community can know p (i.e. the Red Sox won) on the basis of very weak justification simply because no one in *that* community has much stake in whether p ? I argue they are not. Even if I live in a community that does not care much about baseball results, it is plausible that *someone* cares about them. What matters for the communal impurist is whether it is reasonable to expect someone to care about whether p , including people outside of one's local community.

A lot turns on what we mean by a 'community,' so I'll try to be more precise. On my account, an agent knows that p only if he is in a strong enough epistemic position with respect to p to eliminate all of the not- p possibilities that are relevant alternatives to "members of the epistemic community that might draw on the agent's information". As mentioned in Sect. 3, this proposal raises difficult questions about who counts as a member of one's epistemic community. I claimed we should think of the epistemic community as roughly comprised of anyone who might actually draw on one's information, where "might" tracks the notion of a possibility that could reasonably be expected to occur. On this view, an 'epistemic community' is much broader than, say, a group of people living in roughly the same area (i.e. a tribe, village, town, etc.). It is also broader than a group of people who share a language, a religion, and so forth. Although there are varieties of social, geographical, linguistic, and religious communities (among others), I argue there is a sense in which there is only one epistemic community. Let me elaborate.

As inquirers, we all belong to a community of inquiry—a collection of people motivated by the shared need for reliable information to guide our actions and satisfy our natural curiosity. Our reflective understanding of knowledge is that it is based on needs so central to human life that it can be abstractly modeled in a way that, while it might be culturally elaborated in particular ways, is culturally invariant.²⁷ The human interests that ground our concept of knowledge are so basic and so important that this conclusion is inescapable. This is why we have epistemic standards that must be met not only by our linguistic, social, or religious communities, but also by our epistemic community, which is far more expansive.²⁸ We all have a need for reliable information, yet human interests are plastic and unpredictable, so we require informants to meet a sufficiently high standard before

²⁷ See Hannon (2015) for a detailed defense of this view.

²⁸ Given how expansive the epistemic community is, one might prefer a label other than 'community'. You might say we are all members of the same epistemic world [see Thomas (2008) for a similar idea].

we certify their information to others. Even if no one in my town cares about baseball results, there are certainly members of my ‘epistemic community’ who might reasonably care. Thus, our standards for how reliable an informant must be to count as a knower will be sensitive to a diverse range of interests and concerns, including those of people outside my linguistic, social, and religious communities. This explains why communal impurists are not committed to the view that it is easier for, say, the members of a certain village to know that p simply because no one in that ‘community’ has much stake in whether p .

However, appealing to the existence of an epistemic community does not fully resolve the worry. Even if we expand our notion of a community to include the vast majority of people, it is still plausible that certain truths are so trivial that *nobody* would reasonably care about them. Thus, we still encounter *The Problem of Communally Irrelevant Propositions*.

Grimm (2015) attempts to solve this problem. He argues that because humans have diverse practical interests, there might be people with an interest in *whatever* random topic is at issue. If this claim were correct, it would allow us to solve both *The Problem of Communally Low Stakes* and *The Problem of Communally Irrelevant Propositions*. Unfortunately, this idea is implausible. Grimm writes,

I think the right answer is that, given how plastic and unpredictable our practical concerns can be, there is always a story one might tell about why a topic might be of interest to someone, no matter how trivial or insignificant it might seem on the whole. Thus even if getting to the truth about the 323rd number of the Wichita, Kansas phone directory might be (and presumably is) something that you and I could not care less about, it is easy enough to imagine someone who might care about this, if only because he or she wants to phone the person up. More generally, the very idea that there are certain topics that are necessarily trivial or unimportant, from a practical point of view, as opposed to just contingently trivial or unimportant, seems like a mistake. (Grimm 2015: 127–8)

What Grimm says in this passage is only partially right. Our practical concerns are indeed plastic and unpredictable, so a vast range of seemingly trivial truths might nevertheless be important to someone.²⁹ And if our judgments about knowledge are sensitive to the practical concerns of others, this will push the threshold for knowledge high enough to respect these concerns. But even if we are sensitive to the plasticity and unpredictability of peoples’ practical concerns, there is not always a (reasonable) story to tell about why a trivial fact might be of interest to someone.³⁰

²⁹ That said, Grimm’s example of the 323rd number in the Wichita phone directory is problematic. Suppose the 323rd number is Kirstie Alley’s. While it is plausible that someone might want to look up her number, it is unlikely that anyone will want to look up the 323rd number (without caring whose number it might be). Grimm interprets this case in the former way, but I think the latter interpretation is the one Alvin Goldman intended when he originally presented this case. On the latter reading, this example doesn’t support Grimm’s argument.

³⁰ In the passage quoted above, Grimm does not explicitly say that the story one might tell must be *reasonable*, but he is committed to this idea. As mentioned in Sect. 3, Grimm claims the threshold for

Take my earlier example: it is doubtful that anyone does or would care about the number of grains of sand currently on Waikiki Beach. Contrary to Grimm's point, it is not easy to imagine someone who might (in the relevant sense) care about this fact. Rather, it seems like a fact in which *nobody* will likely have any stake.

Thus, it is implausible that for *any* random topic at issue, there might reasonably be someone with an interest in that topic. There are some topics whose truth nobody could reasonably care about. So even though our judgments about knowledge are sensitive to the practical concerns of others, the stakes of some third party will not push the threshold for knowledge to a sufficiently high level in the case of many trivial propositions. For this reason, appealing to the plasticity and unpredictability of our practical concerns will not fully answer these objections.

I want to propose a different solution, one that will allow us to address *The Problem of Communally Low Stakes* and *The Problem of Communally Irrelevant Propositions* in the same way. This solution appeals to the intuitive idea that 'like cases should be treated alike'. Although this principle is more commonly invoked in ethical and legal issues, it also bears on justification and knowledge. To clarify this idea, I'll start with an example.

I have absolutely no interest in the number of active artificial satellites currently orbiting Mars. Nevertheless, it is reasonable to suppose that someone does (or will) have an interest in this question, such as a NASA scientist. Thus, to know the number of active artificial satellites currently orbiting Mars, I must meet a sufficiently high standard. Now let's also suppose, perhaps somewhat unrealistically, that we cannot imagine anyone reasonably caring about the number of *inactive* artificial satellites orbiting Mars (assume NASA has use only for active satellites). If we cannot reasonably imagine anyone caring about these inactive satellites, then it seems to follow (on the impurist view) that it is easier to know the number of inactive satellites orbiting Mars than to know the number of active satellites orbiting Mars. But this is an absurd result. Whatever level of justification is required to know the number of active satellites orbiting Mars, a similar level of justification is presumably required to know the number of inactive satellites orbiting Mars. How can we explain this?

A reasonable explanation is that we must treat like cases alike. Thus, if a certain level of, say, perceptual justification and technological sophistication is required to know the number of active satellites orbiting Mars, then a similar level of perceptual justification and technological sophistication will, *ceteris paribus*, be required to know the number of inactive satellites orbiting Mars. Put more generally, if a certain level of justification (whether perceptual, inferential, mathematical, etc., or some combination) is required to know that *p*, then a similar level of justification will be required to know similar truths.³¹ Further, this principle seems to hold for the various ways of knowing, such as testimony, memory, and reason (inference).

Footnote 30 continued

knowledge will reach a level high enough to respect the "typical" or "normal" stakes of people in the epistemic community (and not stakes that are merely logically possible but unlikely or unreasonable).

³¹ This is similar to a standard way of arguing in ethics, namely, if we should do *X* in circumstances *C*₁, and circumstances *C*₂ are just like *C*₁, then we should do *X* in *C*₂.

Consider the following case of testimonial knowledge. Suppose I want to know whether the phrase “There’s daggers in men’s smiles” is from Hamlet or Macbeth. Let’s also stipulate that I have no Internet access, so I am unable to look up the answer at home. Fortunately, my neighbor, Jess, is a renowned expert on the works of William Shakespeare. I decide to ask Jess and she tells me with unwavering confidence that the phrase appears in Act II, Scene III of Macbeth. Jess is right, and I unhesitatingly form the corresponding true belief.

It should provoke no controversy to say this is a case of testimonial knowledge. As an expert on the subject, Jess clearly meets the level of justification required to know the answer to my question. Further, I can acquire knowledge by trusting her testimony.

Now let’s modify the example. It is easy to imagine why someone might (in the relevant sense) care about the correct reference for a famous phrase like “There’s daggers in men’s smiles,” but it’s difficult to imagine why someone might care about the correct reference for an inconsequential phrase like “I thank you, doctor” (Macbeth Act IV, Scene III). This particular phrase seems so unimportant that it is unlikely anyone will care much, if at all, about it. But if we cannot reasonably imagine someone caring about the source of this phrase, then we cannot fix the threshold for knowing this truth by appealing to its importance to some possible actions. This leads us back into *The Problem of Communally Low Stakes* and *The Problem of Communally Irrelevant Propositions*. If we cannot imagine someone having much stake in this issue, then communal impurism seems to imply that it is more difficult to know that “There’s daggers in men’s smiles” is from Macbeth than to know that “I thank you, doctor” is from Macbeth. And if we cannot imagine anyone having *any* stake in this issue, then communal impurism will fail to fix the level of justification required to know this truth. As a result, it will leave knowledge of this truth indeterminate. Both consequences are theoretical costs.

To fully resolve this issue, we must explain why good evidence is required not only to know truths that might reasonably matter to someone, but also why good evidence is required to know truths that nobody has much, if any, stake in. As mentioned above, an intuitively plausible explanation is that we treat like cases alike.³² Thus, whatever level of justification is required to know that one quote is from Macbeth, a similar level of justification will, *ceteris paribus*, be required to know that another quote is from Macbeth. More generally, if a level of justification n is required to know that p , then a similar level of justification is required to know truths that are sufficiently similar to p .

What counts as a ‘sufficiently similar’ truth (or which cases qualify as ‘alike’) is a difficult conceptual issue. I doubt we can articulate a set of general criteria to draw this distinction in a principled way. Nevertheless, it is reasonable to suppose that our judgments about these cases will (or would) sufficiently coincide for us to say tolerably well, in particular cases, which truths require similar levels of justification. This is illustrated by the two examples discussed above. We intuitively recognize

³² The policy of treating like cases alike might threaten to leach to absolutely all cases via the transitivity of “x-is-similar-to-y.” However, we can resist this by proposing a two-tiered policy: only add those cases that are similar to a base case (i.e. to a case captured by the original communal impurist principles).

that if a certain threshold must be met to know the number of active satellites orbiting Mars, then, *ceteris paribus*, a similar threshold must be met to know the number of inactive satellites orbiting Mars. In other words, we recognize these cases are sufficiently similar to demand similar thresholds for knowledge. This is also true for the case of testimonial knowledge mentioned above.

Why do we treat like cases alike? Is there some deeper explanation for why this principle seems to inform our judgments about knowledge? I'll conclude by sketching one possible answer, namely, that using comparable epistemic standards has some utility, whereas applying widely divergent thresholds in similar contexts would be a confusing and inefficient epistemic strategy.

To illustrate this, let's revisit the sand on Waikiki Beach. As I've argued, it is doubtful that anyone does (or would) care about the number of grains of sand currently on Waikiki Beach. I, for one, cannot imagine anyone reasonably appealing to the proposition "Waikiki Beach has n grains of sand" in their practical reasoning. However, it is not unreasonable to imagine an arenophile (someone who collects sand as a hobby) who might care about counting grains of sand in *certain* contexts. For example, Dilbert the arenophile might care that he has found n grains of gypsum in a small heap of sand, since gypsum is a rare type of sand. In order to know he has n grains of gypsum, Dilbert must therefore count them carefully. Put differently, he must meet a certain minimum threshold to know there are n grains of sand. Further, our ability to imagine someone with Dilbert's practical interests explains why, on the communal impurist view, one must satisfy a sufficiently high minimum threshold to know there are n grains of gypsum. And if the thresholds for knowledge are fairly standardized across contexts (bracketing cases in which someone has especially high stakes), then we should expect agents to be in a similarly strong epistemic position to know, for any similarly sized heap of sand, there are n grains in it.³³

But *why* expect the thresholds to be symmetrical in this way? We gain clarity on this question by considering an alternative view, namely, the possibility of widely diverging standards for knowledge in similar cases. When we reflect on this possibility, we can recognize how epistemically inefficient such a practice would be. First, it introduces opportunities for miscommunication because, on this view, we would be relying on each individual to judge whether or not a given proposition is one in which an inquirer might reasonably have stake, but people's judgments about this might not coincide. As a result, there will be diverging thresholds for knowledge, which makes epistemic evaluation less effective. Second, the use of widely diverging standards in similar cases places additional cognitive burdens on each member of the epistemic community, since they must keep track of the various thresholds for knowledge. But this might stretch our cognitive load beyond its capacities. A community-wide system of testimony that worked this way would be inefficient because widely diverging standards would make it more difficult to pool and share potentially useful information, and informational exchange might get

³³ The heap of sand must be 'similarly sized' because it becomes more difficult to count grains of sand as the number goes up. When counting, I would be more certain there are three grains before me than 200.

bogged down. Third, there seems to be something deeply inconsistent about an epistemic system that allows two informants, in roughly similar situations, to both merit the honorific label ‘knower,’ even though one of them is in a strong epistemic position and the other is in a very weak epistemic position. If this account were correct, we would demand far more justification to know the number of active satellites orbiting Mars than to know the number of inactive satellites in Mars’ orbit. Likewise, we would demand more or less expertise to locate a specific passage in Shakespeare depending on whether the passage is one we could reasonably imagine someone caring about. But an epistemic system that permitted such differing thresholds would strike us as inconsistent.

In contrast, by demanding similar thresholds in similar cases, we achieve a consistency that promotes a deep kind of coordination in our basic epistemic practices.³⁴ It allows for the coordination of epistemic rule following across the community, which in turn makes testimony more trustworthy and reliable. In this way, our demand for similar knowledge thresholds is importantly connected to the communal view of knowledge. The practice of having similar thresholds gives rise to a more efficient epistemic economy because the use of similar knowledge thresholds allows us to each recognize that we are following the same rules and standards. As a result, we don’t have to worry about people applying diverging thresholds in similar cases. When consistent thresholds are shared, I can trust that *you will* draw the same conclusion from an evidential basis that *I would*. For these reasons, we typically demand similar levels of justification for similar cases, regardless of whether we can reasonably imagine someone with practical stake in the proposition.

Brown worries that setting a minimum standard for knowledge potentially gives rise to the problem of arbitrariness (2014: 187). What would fix the minimum standard for knowledge? Without a plausible answer, we land right back into the threshold problem, which is the problem of how to provide an account of what fixes the level of justification required for knowledge in a non-arbitrary way.

But the problem of arbitrariness does not arise for the communal approach. The minimum standard for knowledge derives from our need to identify good sources of information to members of our community. Basic practical concerns thus generate a standard that is fitting to certify actionable information for many people with diverse interests. This sets a general threshold for knowledge in a non-arbitrary way. The level of justification needed for knowledge is that which puts the agent in a strong enough epistemic position for her to serve as a reliable source of actionable information for her practical reasoning environment, and minimally for the normal standards of her epistemic community.

Brown might also reject communal impurism on the grounds that “appeal to imaginary scenarios does not properly restrain the requirements for knowledge” (2014: 188). Sticking with her Mars example, Brown worries that we can imagine a scenario in which a lot rests on the proposition that Mars has two moons as well as a scenario in which very little rests on that proposition. This is supposed to illustrate

³⁴ See Dogramaci (2012) for a similar approach to rationality.

that “appeal to merely imaginary scenarios does not appear to give us a determinate answer to the threshold problem since we can imagine any scenario we like with any stakes whatsoever” (2014: 188).

According to the communal approach, however, the threshold for knowing a proposition that is irrelevant to any practical reasoning situation faced by the subject is not set by whatever scenario the subject happens to imagine. Rather, that threshold is set by the possibilities that are relevant alternatives to members of the epistemic community that might draw on the information. Of course we cannot anticipate every possible scenario in which such information might be put to use, nor must we imagine all of these scenarios when ascribing knowledge in daily life. But, as I've argued, we do have a rough idea of what normally counts as having done enough to establish the propriety of a knowledge claim. As information-dependent and information-sharing creatures, we are accustomed to the fact that others often turn to us for information about various topics. But since it is hard to say in advance who will come to rely on our judgments, the threshold for knowledge gravitates towards a level high enough to respect the “typical” or “normal” stakes of others who might appeal to these judgments. This is the communal threshold for knowledge, but this standard can be overridden when the stakes (whether one's own or someone else's) in the relevant practical reasoning environment are especially high.

6 Against subject-centered impurism

I have argued that contrary to what many impurists believe, it is not necessarily the practical reasoning situation faced by the *subject* at a time that determines whether she knows. If a knower is someone on whom many people can rely, it makes little sense to think that the epistemic standards would be so closely tied to the interests of the subject. We are interested in whether an informant is epistemically positioned to provide information that is fitting to the community's interests, not just the subject's interests. Thus, the threshold for knowledge will be strongly conditioned by the diverse needs and interests of the community to which the subject belongs. Having such a standard is what makes pooling and sharing information easier. Thus, we have a reason to prefer community-centered impurism to subject-centered impurism.³⁵

I agree with subject-centered impurists that the threshold for knowledge can be overridden when the stakes for the subject are especially high. However, my view does not collapse into subject-centered impurism for two reasons. First, the practical interests of the subject do not necessarily set the standard for knowledge. I have argued that whether someone knows is partly determined by the interests of other members of the community; for instance, I should not recommend to you an informant who barely meets the communal standard if I am aware that your stakes are exceptionally high. Second, I have argued that the subject's stakes cannot drag

³⁵ Greco (2008) and Henderson (2009) raise a similar objection to impurism.

the threshold for knowledge lower than the communal level (and I have explained why this asymmetry is not problematic). In contrast, subject-centered impurists claim that a person in a relatively low-stakes situation counts as knowing that p if that person is able to rule out the not- p possibilities from a range of alternatives that is narrower than the range demanded by the community.³⁶ In these ways, the communal view is importantly different from subject-centered impurism.

7 Conclusion

Standard formulations of impurism are vulnerable to serious objections, which casts doubt on the impurist's ability to solve the threshold problem. The root of the problem for impurism has been the idea that the level of justification required for knowledge is set by the *subject's* practical reasoning situation. However, I have outlined a new way of implementing the impurist strategy that jettisons this idea in favor of a more plausible view: the threshold for knowledge is set by a communal standard, and the required level of justification can increase if the relevant practical reasoning situation calls for it. By adopting this view, we end up with a plausible version of impurism that answers the threshold problem and is not vulnerable to Brown's criticisms against the unity and relevance approaches.

Acknowledgments Thanks to Nathan Ballantyne, Sinan Dogramaci, Elizabeth Edenberg, Mikkel Gerken, Stephen Grimm, Nick Hughes, Andy Mueller, Coran Stewart, and Emily Sullivan for comments on an earlier draft. This paper was presented at the Young Philosophers Lecture Series, and I am grateful to the audience there for feedback.

References

- Austin, J. L. (1946). Other minds. *Proceedings of the Aristotelian Society*, 20, 148–187.
- BonJour, L. (2010). The Myth of knowledge. *Philosophical Perspectives*, 24(1), 57–83.
- Boyd, K., & Nagel, J. (2014). The reliability of epistemic intuitions. In E. Machery & E. O'Neill (Eds.), *Current controversies in experimental philosophy*. Routledge.
- Brown, J. (2006). Contextualism and warranted assertibility manoeuvres. *Philosophical Studies*, 130(3), 407–435.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Brown, J. (2014). Impurism, practical reasoning, and the threshold problem. *Noûs*, 48(1), 179–192.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Oxford University Press.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- Dogramaci, S. (2012). Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research*, 84(3), 513–530.

³⁶ It is for similar reasons that I reject Henderson's view (2009). Although we agree that knowledge ascriptions certify people as reliable informants to members of their epistemic community, Henderson maintains that "the community in question may be [...] composed of just the attributor and the interlocutor" (2009: 123). If this were correct, however, and if neither the informant nor her interlocutor had practical stake in the matter, then *The Problem of Easy Knowledge* would easily resurface. Thus, an implausible consequence of Henderson's view is that the standard for knowledge may be set *too low*.

- Dogramaci, S. (2015). Forget and forgive: A practical approach to forgotten evidence. *Ergo*, 2(26), 645–677.
- Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40(3), 363–378.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1), 67–94.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Fricke, M. (2008). Skepticism and the genealogy of knowledge: Situating epistemology in time. *Philosophical Papers*, 37(1), 27–50.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178(3), 529–547.
- Gerken, M. (2015). The roles of knowledge ascriptions in epistemic assessment. *European Journal of Philosophy*, 23(1), 141–161.
- Greco, J. (2008). What's wrong with contextualism? *Philosophical Quarterly*, 232, 416–436.
- Grimm, S. (2011). On intellectualism in epistemology. *Mind*, 120(479), 705–733.
- Grimm, S. (2015). Knowledge, practical interests, and rising tides. In J. Greco & D. Henderson (Eds.), *Epistemic evaluation: Point and purpose in epistemology*. Oxford: Oxford University Press.
- Hannon, M. (2013). The practical origins of epistemic contextualism. *Erkenntnis*, 78(4), 899–919.
- Hannon, M. (2014). Fallibilism and the value of knowledge. *Synthese*, 191, 1119–1146.
- Hannon, M. (2015). The universal core of knowledge. *Synthese*, 192(3), 769–786.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Henderson, D. (2009). Motivated contextualism. *Philosophical Studies*, 142(1), 119–131.
- Hetherington, S. (2001). *Good knowledge, bad knowledge: On two dogmas of epistemology*. Oxford: Oxford University Press.
- Hetherington, S. (2006). Knowledge's boundary problem. *Synthese*, 150, 41–56.
- Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., et al. (2015). Gettier across cultures. *Noûs*, 50(2). doi:10.1111/nous.12110.
- McKenna, R. (2013). 'Knowledge' ascriptions, social roles, and semantics. *Episteme*, 10(4), 335–350.
- Nagel, J., San Juan, V., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129, 652–661.
- Neta, R. (2006). Epistemology factualized: New contractarian foundations for epistemology. *Synthese*, 150(2), 247–280.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3), 247–279.
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107(2), 143–157.
- Russell, G. K., & Doris, J. M. (2008). Knowledge by indifference. *Australasian Journal of Philosophy*, 86(3), 429–437.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stanley, J., & Hawthorne, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Thomas, A. (2008). The genealogy of epistemic virtue concepts. *Philosophical Papers*, 37(3), 345–369.
- Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1–2), 429–460.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Turri, J. (2013). A conspicuous art: Putting gettier to the test. *Philosophers' Imprint*, 13(10), 1–16.