

# The universal core of knowledge

Michael Hannon

Received: 24 July 2014 / Accepted: 23 October 2014 / Published online: 21 November 2014  
© Springer Science+Business Media Dordrecht 2014

**Abstract** Many epistemologists think we can derive important theoretical insights by investigating the English word ‘know’ or the concept it expresses. However, native English speakers make up less than 6 % of the world’s population, and some empirical evidence suggests that the concept of knowledge is culturally relative. So why should we think that facts about the word ‘know’ or the concept it expresses have important ramifications for epistemology? This paper argues that the concept of knowledge is universal: it is expressed by some word in every natural language. I also explore the implications of this thesis for philosophical methodology.

**Keywords** Knowledge · Experimental philosophy · Philosophical methodology · Epistemic relativism · Practical explication

## 1 Introduction

A deep presupposition guiding mainstream epistemology is that we can derive important theoretical insights by investigating the English word ‘know’, the concept it expresses, and/or the semantic properties of sentences of the form “S knows that *p*”. This presupposition is manifest in a number of diverse strands of contemporary research; for example, folk attributions of knowledge have been used to support (and to challenge) epistemic contextualism (DeRose 1995), subject-sensitive invariantism (Stanley 2005), the knowledge norm of assertion (Williamson 2000), the knowledge norm of practical reasoning (Hawthorne and Stanley 2008), and group knowledge (Lackey 2012).

---

M. Hannon (✉)

Department of Philosophy, Fordham University, 441 E. Fordham Road, Bronx, NY 10458, USA  
e-mail: mhannon09@gmail.com

However, native English speakers make up less than 6% of the world's population. Though English is the world's third most common language, it is just one of over 6,000 spoken languages. So why should we think that facts about the English word 'know' have important ramifications for epistemology?

One significant worry is that epistemologists are only doing so-called 'ethno-epistemology' (Nichols et al. 2003; Weinberg et al. 2001, reprinted 2008). An 'ethno-epistemologist' describes the epistemic terms, concepts, norms, or practices of some specific group of people. While there is nothing wrong with this research program per se, most epistemologists take themselves to be investigating *general* epistemic facts. For example, skepticism is worrying and interesting presumably because it purports to show that there is no human knowledge (or almost none). Similarly, the philosophical significance of epistemic contextualism would be diminished if the context-sensitivity of 'know' were peculiar to English. Epistemological investigation typically aspires to general conclusions, not culturally local ones.

This raises important methodological questions about the nature and aims of epistemological theorizing. Should epistemologists be interested in the concept expressed by the English word 'know'? Can language tell us anything of significance to epistemology?

I will argue that important epistemological insights can be derived by investigating the word 'know', the concept it expresses, and the semantic properties of "S knows that *p*". Section 2 will briefly discuss whether the concept of knowledge is 'universal' in the sense that this concept is expressed by some word in every (or virtually every) natural language. Section 3 will provide a theoretical argument for why language users living in social communities would develop a concept of knowledge that is shared cross-culturally. Section 4 will discuss the objection that the concept expressed by the English word 'know' is *not* universal. Section 5 will argue that the concept of knowledge has a sufficiently large universal core to say that this concept may be found across cultures. Section 6 explores the implications of this thesis for philosophical inquiry.

By defending the thesis that the concept of knowledge is universal, I will show that epistemologists who study facts about the use of English word 'know' are not just practicing so-called 'ethno-epistemology'. However, I will also identify some philosophical projects that *are* vulnerable to the charge of ethno-epistemology. This illustrates when and why epistemologists may appeal to facts about the English word 'know' for theorizing.

## 2 Is the concept of knowledge universal?

According to the program of semantic research known as Natural Semantic Metalanguage, there exists a small inventory of universal concepts that are expressed by some word in every human language (Goddard and Wierzbicka 2002). Evidence from cross-cultural semantics suggests that the concept expressed by the word 'know' is universal in this sense (Goddard 2010).<sup>1</sup> Moreover, Haspelmath and Tadmor (2009) report that a

<sup>1</sup> More precisely, the sense of 'know' that embeds a propositional or wh- complement is arguably universal. Not every language follows English by making 'know' perform multiple tasks: French has

word for ‘know’—in the sense that embeds a propositional complement—features in all the languages in the World Loanword Database, which includes a broad sampling of languages from every inhabited continent. Thus, linguists have isolated ‘know’ as one of a few words that are universal.

If ‘know’ were universal in this way, it would be a remarkable fact because empirical evidence from cross-linguistic semantics indicates that almost every word in the English lexicon does *not* have an equivalent in many, perhaps most, other languages. Even the meanings of words for seemingly universal epistemic states and cognitive processes like ‘believe’, ‘doubt’, and ‘remember’, as well as words for emotional states like ‘sad’ and ‘angry’, are language and culture specific (Russell 1991; Wierzbicka 1999; Amberber 2007; Goddard 2010). The universality of the word ‘know’ would thus cry out for an explanation. Why is it that no human language seems to get by without a word that accurately translates ‘know’?

Following a suggestion made by Craig (1990), I hypothesize that the concept of knowledge is universal because it fulfills important needs and serves essential functions in human life and thought. More specifically, I will defend the following proposal about the point or purpose of the concept of knowledge: we have a concept of knowledge to *identify sufficiently reliable informants*.<sup>2</sup> This puts the main idea very briefly and roughly (i.e. reliable for who, and what counts as “sufficient”?); I will develop it in Sect. 3. I will show that understanding the social role of the concept of knowledge will contribute to understanding why every human society can be expected to possess this concept.<sup>3</sup> My paper not only provides theoretical support for the view that ‘know’ is a linguistic universal but also goes beyond the existing research by explaining why every sufficiently developed language-using society will have the concept of knowledge.

The Natural Semantic Metalanguage (NSM) program has much relevance to epistemology. If the data are reliable, then epistemologists studying the word ‘know’ are not merely investigating the English verb but rather a concept that we all have in spite of the language(s) we happened to have learned as children. This would imply that by studying the concept expressed by the English word ‘know’ we could learn something about the universal concept of knowledge. The claim that epistemologists cannot derive important philosophical insights by studying the English word ‘know’ would therefore misfire.

---

Footnote 1 continued

‘savoir’/‘connaitre’ and German has ‘wissen’/‘kennen’. Evidence suggests that the situation is more complicated *vis-à-vis* translations of ‘know how’. While it might be true that propositional and question-embedding verbs are not different across languages, this is clearly not true for all uses of ‘know’. My focus is on propositional knowledge, and I remain neutral on whether knowledge-wh reduces to propositional knowledge.

<sup>2</sup> Many people doubt that the concept of knowledge has just one purpose (e.g. Rysiew 2012; Beebe 2012). In this paper I merely assume that flagging reliable informants is a central (common, important) function of the concept of knowledge, while remaining open to the possibility that we may need to supplement or modify this hypothesis if it does not do all the work we need. Pluralism about the functional roles of the concept of knowledge seems highly plausible.

<sup>3</sup> I do not claim that every society *must* possess the concept of knowledge, but I will argue that a society must have this concept if they are to meet certain basic pragmatic needs.

While NSM research provides some evidence that the concept of knowledge is universal, I will not critically engage with this literature. My approach in this paper will be that of an armchair philosopher. I will leave open the empirical question of whether the linguistic data supports the thesis that the concept expressed by the word ‘know’ is universal. My aim is to supplement this research by providing a theoretical argument for why this concept is universal. More specifically, I will explain why language users like us, living in social communities, would develop a cross-culturally and cross-linguistically shared concept of knowledge. My argument is that all humans have certain basic practical needs that the concept of knowledge satisfies, and we could not meet these needs with an epistemic lexicon that does not contain any words with a content that is (near-) synonymous with ‘knowledge’; thus, we should expect every culture to have a word to express the concept of knowledge.<sup>4</sup>

This theoretical argument can be refuted empirically. Cross-linguistic semantics might eventually reveal that the concept of knowledge is culturally relative in the way that many other concepts are language and culture specific. What if it turns out that the concept expressed by the word ‘know’ is not universal? Section 4 will review some existing data that seems to support this conclusion. This raises a significant worry: what is the philosophical significance of a concept that we acquired as a result of the seemingly arbitrary fact that we happened to learn one language and acquire certain concepts rather than others? As I’ll argue, we should doubt that any sufficiently developed language-using society would lack the concept of knowledge. Human societies tend to adopt (or create<sup>5</sup>) concepts that are useful to them, and any society that lacked the concept of knowledge would be significantly epistemically impoverished. However, I will conclude by outlining a way of doing normative epistemology that is compatible with the possibility that the word best translated as ‘know’ in some (actual or hypothetical) societies *does* have a different meaning from the English word.

### 3 A pragmatic theory of the concept of knowledge

#### 3.1 The method of practical explication

Why do language users like us have a concept of knowledge? To answer this question I will use the philosophical method of practical explication, which was developed by Craig in his 1990 book *Knowledge and the State of Nature*. A practical explication of a concept seeks to elucidate the role (or roles) that the concept plays in our socio-linguistic interactions, given a set of facts about humans, their environment, and their interests or aims. How does one ascertain what role(s) the concept of knowledge plays and what conceptual needs it satisfies? Craig describes this methodology as follows:

<sup>4</sup> A near-synonymous word—‘schnowledge’—could perhaps do the job of ‘knowledge’ by playing the same functional roles. However, it is unlikely that we could meet our present communicative needs with an epistemic lexicon that does not contain any words with a content that is at least nearly synonymous with ‘knowledge’. Thanks to Mikkel Gerken for pointing this out.

<sup>5</sup> I will sometimes speak of “adopting” concepts, but my argument does not depend on anything like Platonism.

We take some *prima facie* plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application. . . And if [there is such a role], one way to find out must be to form some hypothesis about it, try to work out how a concept custom-designed for that role would look, and then see to what extent it matches our everyday practice with the concept of knowledge as actually found. (Craig 1990, pp. 2–3)

Having sketched his method, Craig then states his hypothesis: the purpose of the concept of knowledge is to flag reliable informants (1990, p. 35).

This hypothesis emerges from the idea that the concept of knowledge facilitates the survival and flourishing of human beings in communities. We need true beliefs about our environment in order to successfully guide our actions, so we require sources of information that will lead to true beliefs. Often the easiest and most efficient way to acquire a true belief is to ask somebody reliable; and since on almost any issue some informants will be more likely than others to provide a true belief, any community may be presumed to have an interest in evaluating sources of information. Craig thus takes as his starting point the idea that the concept of knowledge has developed to play an important role in the general economy of our concepts, namely, to certify reliable informants.

On this hypothesis, a species with a radically different psychology from ours might not have anything like the concept of knowledge. For example, a community of omniscient beings would have no need for the concept of knowledge because there would be no need for them to distinguish reliable informants from unreliable informants. It is plausible that our epistemically evaluative concepts arise in connection with certain needs and practices, so concepts that do not serve any function are unlikely to be adopted. In fact, we should expect this reasoning to apply to any word or concept in our linguistic and conceptual repertoires. This point is closely related to Hume's insight that a community of perfectly altruistic beings in a land of abundant resources would have no need for a concept of justice (*Treatise* 3.3.2). We are social animals and we will adopt words and concepts that are useful to us. The more basic the human needs are that the concept we are explicating satisfies, the more essential this concept will be to human life and thought.<sup>6</sup>

The hypothesis I will explore is that the concept of knowledge is used to satisfy certain basic practical needs; in particular, this concept functions to identify the adequacy of informants for people and purposes beyond our own immediate concerns, which allows us to pool and share information. I'll now elaborate on this point.

### 3.2 The importance of the concept of knowledge

We all have a need to find informants who are reliable enough for our own purposes: I want someone reliable enough for *me*, you want someone reliable enough for *you*, and so on. However, this picture makes it very difficult to volunteer oneself or another as a

<sup>6</sup> A consequence of this view is that the concept of knowledge is species relative. Knowledge is something that humans have, but it might not be something that more (or less) intelligent creatures possess.

good informant. We would be required to keep track of each individual's interests and purposes, but we usually do not have such information. Thus, identifying informants who are only reliable enough for our own purposes is inadequate as a general strategy for human epistemic practices. Our needs and practices are more complex than just wanting information for our purposes here and now. We have a need to share information and identify reliable informants *in general*. Since we often do not know the practical needs or interests of persons to whom an informant might be useful, we need a way to ensure that informants are reliable for a wide range of people and purposes.

Craig therefore appeals to a process he calls *objectivization* (1990, p. 82). Jonathan Dancy describes this process in the following way: “A concept is objectivized if it becomes progressively less tied to the particular concerns of the user” (1992, p. 394). We abstract away from details of the particular circumstances surrounding one individual's inquiry to arrive at a more intersubjective, socially-directed concept. We are not just evaluating whether someone is reliable enough for our own individual purposes, but also whether that person is reliable enough for other inquirers. Knowledge attributions certify an informant's belief or information to a shared epistemic community. When we attribute knowledge to someone, we certify that the agent is epistemically positioned such that people can freely draw on his or her information.

The ‘relevant alternatives’ framework provides a way to articulate this idea. What does it take to qualify as knowing that *p*? The 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. Which alternatives are relevant to the members of the agent's epistemic community and which are not? Many think that this question has not been answered with a satisfying degree of precision, which has led critics to view the relevant alternatives theory as *ad hoc* or obscure. However, it is plausible that we have a rough but reliable sense of what counts as having done enough to establish the propriety of a knowledge claim. Certain alternatives are not the kind that we humans take to be the likely counter-possibilities to what the subject is said to know, and thus we need not rule them out.<sup>7</sup> For example, it does not seem that I can be legitimately criticized or blamed for failing to rule out the possibility that I am a brain-in-a-vat when making epistemic evaluations in everyday life, at least not in the same way that I could be criticized for failing to take into account the possibility that (say) the bird in my garden is a goldfinch when claiming to know that it is a canary (Austin 1961, p. 83).

This account illustrates why the concept of knowledge is deeply important to human life.<sup>8</sup> Knowledge attributions certify informants who are in a sufficiently high-quality epistemic position such that inquirers may draw on their information for a wide variety of projects and purposes. That this concept arises from the most basic needs to identify,

<sup>7</sup> I do not expect there to be unanimous agreement about all of the possibilities that need to be eliminated in order to qualify as knowing. I merely presume that such judgments will (or would) coincide sufficiently to give us what Patrick Rysiew calls “a set of ‘core’ not-*p* alternatives” (2001, p. 489). Without this assumption it would render utterly mysterious how people are adept at determining what a speaker means in uttering “S knows that *p*” (*ibid.*).

<sup>8</sup> I further develop this view in Hannon (2014).

pool, and share reliable information also explains why we should expect this concept to arise in any human society. My argument is not just that the concept of knowledge is beneficial in certain ways, but that this concept has the very shape it would have if it were designed with these benefits in view (Gerken 2013 also makes this point).

This contrasts with the usual emphasis that NSM research places on the contingency of our concepts. Linguists such as Wierzbicka (1999) and Goddard (2010) argue that the language of mental state concepts displays a great deal of cultural variation, which reminds us of the contingency of our concept possession and how easily things might have been different. In contrast, my account illustrates that the concept of knowledge is not highly contingent. Given some basic facts about humans, our conditions of life, and some of our interests and aims, we have needs that the concept of knowledge arose to satisfy. This concept does not just happen to satisfy these basic needs; the concept must be this way (or sufficiently similar) in order to satisfy them. Knowledge-talk is a necessary response to our need to identify, pool, and transmit reliable information.

I do not mean to suggest that the concept of knowledge arises out of anything like metaphysical necessity. We need to carefully distinguish between two possibilities: first, it might be that a concept which arose to meet certain needs must, necessarily, have a certain shape in order to meet those needs; second, it might be that it is necessary for us to have a specific concept that meets these needs. I defend the former claim but not the latter. The shape of the concept of knowledge is necessitated by it arising to meet certain needs; moreover, the basic nature of these needs makes it incredibly likely (although not necessary) that we would invent (or adopt) this concept. Whatever the actual history of the concept of knowledge might be, it is required that we have such a concept in order to satisfy certain basic practical needs.

All societies will need to solve the problem of how to identify, pool, and share reliable information with a community of inquirers, so we should expect every sufficiently developed society to possess a concept that does this work.<sup>9</sup> This argument provides support for the idea that ‘know’ is a linguistic universal and vindicates the common methodological view according to which we can derive important theoretical insights by investigating facts about the word ‘know’. If the concept expressed by the word ‘know’ is universal, then studying the word ‘know’ (and sentences of the form “S knows that *p*”) will shed light on the universal concept of knowledge.

#### 4 Epistemic relativism

What if it turns out that the concept expressed by the English word ‘know’ is not universal? The claim that ‘know’ is a linguistic universal is not uncontroversial; some empirical evidence indicates that the concept of knowledge is culturally relative (see Nichols et al. 2003; Weinberg et al. 2001). I will discuss this research in what follows.<sup>10</sup>

<sup>9</sup> At some point in human history societies might not have been far enough along in their development to effectively address these needs.

<sup>10</sup> Weinberg et al. (2001) also report that epistemic judgments differ based on socio-economic status and educational background. My paper focuses specifically on cultural differences, but most of what I say applies equally to these other factors. Some people have claimed that there are gender

A concept can be culturally relative in at least two ways. First, there might be systematic variation in the extension of the word ‘know’ as used by English speakers from different cultural backgrounds. For example, people of Western descent might refrain from attributing knowledge to Gettier victims whereas people of East Asian descent might be willing to ascribe knowledge in such cases. This may suggest that they have different concepts. Second, the concept of knowledge might be culturally relative because the word best translated as ‘know’ in some (actual or hypothetical) society expresses a different concept from that expressed by the English word ‘know’. How worrying are these possibilities of conceptual variation?<sup>11</sup>

#### 4.1 Some evidence that epistemic judgments are culturally relative

In a widely cited study, [Weinberg et al. \(2001\)](#) report cultural variation in epistemic judgments. To take one example from their study, they ask people to consider the following Gettier case:

*Gettier Case.* Bob has a friend Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it? ([Weinberg et al. 2001](#), p. 29)

Weinberg, Nichols, and Stich (WNS) report that the majority of East Asian (56 %) and South Asian (61 %) participants responded with the judgment that Bob has knowledge (“really knows”), whereas the majority of Western participants (74 %) maintained that Bob lacks knowledge (“only believes”). Thus, Westerners seem more likely to deny knowledge to the victim of this Gettier scenario than their Eastern counterparts ([Weinberg et al. 2001](#), pp. 30–31).

WNS also report that participants of East Asian heritage were less likely than Western participants to ascribe knowledge to the protagonist of Keith Lehrer’s Truetemp case ([Lehrer 1990](#), pp. 163–164). They test the following example:

*Truetemp Case.* One day Charles is suddenly knocked out by a falling rock, and his brain becomes re-wired so that he is always absolutely right whenever he estimates the temperature where he is. Charles is completely unaware that his brain has been altered in this way. A few weeks later, this brain re-wiring leads him to believe that it is 71 degrees in his room. Apart from his estimation, he has

---

Footnote 10 continued

differences between epistemic intuitions ([Buckwalter and Stich 2011](#)), but [Boyd and Nagel \(2014\)](#) and [Starmans and Friedman \(2012\)](#) report no gender differences in epistemic judgments.

<sup>11</sup> I am assuming that epistemic judgments (or “intuitions”) are a basic source of evidence concerning our epistemic concepts (see [Jackson 2011](#) for a similar view). It is my concept of knowledge that enables me to say, for instance, that a Gettier victim does not know, and this reveals something about the shape of that concept. Of course there can be conditions under which our judgments mislead us, just as perceptual seemings are usually reliable but not infallible. Thus, not every epistemic judgment will deliver direct evidence about the epistemic concept in question.



no other reasons to think that it is 71 degrees. In fact, it is at that time 71 degrees in his room. Does Charles really know that it was 71 degrees in the room, or does he only believe it? (Weinberg et al. 2001, p. 26)

While both East Asian and Western subjects were more likely to classify Charles' belief as not a case of knowledge, WNS found that East Asian subjects were significantly more likely to deny knowledge (88 %) than Western subjects (68 %).

Another test by WNS used a version of Fred Dretske's (1970) Zebra case:

*Zebra Case.* Mike is a young man visiting the zoo with his son, and when they come to the zebra cage, Mike points to the animal and says, "that's a zebra." Mike is right—it is a zebra. However, as the older people in his community know, there are lots of ways that people can be tricked into believing things that aren't true. Indeed, the older people in the community know that it's possible that zoo authorities could cleverly disguise mules to look just like zebras, and people viewing the animals would not be able to tell the difference. If the animal that Mike called a zebra had really been such a cleverly painted mule, Mike still would have thought that it was a zebra. Does Mike really know that the animal is a zebra, or does he only believe that it is? (Weinberg et al. 2001, p. 32)

WNS found that only 32 % of Westerners maintained that Mike really knows that the animal is a zebra, whereas 50 % of participants from the Indian Subcontinent attributed knowledge to Mike. This data shows that our knowledge-ascribing judgments sometimes display cultural variation.

If these sorts of conceptual differences exist, we should not be totally surprised. These differences are probably less rare than one might think. Conceptual differences can go unnoticed for decades or whole lifetimes, especially when the difference is slight and the sort of hypothetical cases that would bring it out are unusual.<sup>12</sup>

In epistemology, however, there are substantive theoretical issues at stake. As mentioned at the start of this paper, much epistemological theorizing has been guided by the presupposition that we can derive important theoretical insights by investigating how competent speakers use the word 'know'. Thus, if it turns out that epistemologists have only been investigating a word or a concept that some speakers (e.g., Caucasian English speakers) possess, this would pose a *prima facie* challenge to much epistemological theory. Why should we care about features that are peculiar to the English verb 'know'? What if epistemic contextualism is true for English knowledge attributions but not Japanese? Would the Gettier problem still be significant if most non-Westerners attributed knowledge to the victim of a Gettier scenario? Isn't it arbitrary to privilege the epistemic judgments of one cultural group over another?

The cultural variation cited by WNS challenges the idea that the concept of knowledge is universal. If Westerners and non-Westerners do not share a concept of knowledge, then it cannot be true that the concept expressed by the word 'know' is universal. The upshot is that epistemologists studying the concept expressed by the English word

<sup>12</sup> Lycan (2006) has a nice example involving Sartre's claim that a 'lie' cannot be true. According to Lycan, many people think that someone who intends to provide false information but mistakenly provides true information is still a liar.

‘know’ would be investigating the concept possessed by a particular cultural group, not a concept that we all have despite our different heritages. Thus, the cultural relativity of the concept of knowledge threatens to undermine many insights that have been derived from standard epistemological methodology.

I will provide a twofold response to this challenge. First, I will draw on new empirical research to argue that the data presented by WNS are unreliable (Sect. 4.2). Second, I will argue that even if there were evidence of cultural variation, this need not worry epistemologists (Sects. 5 and 6).

#### 4.2 The WNS findings are not robust

It is doubtful that the findings reported by WNS are reliable. [Boyd and Nagel \(2014\)](#) found no robust evidence that the epistemic judgments of different demographic groups are deeply at odds with each other. Similarly, [Nagel et al. \(2013\)](#) were unable to replicate WNS’s Gettier case results, finding no statistically significant difference in knowledge ascriptions across ethnicity, gender, age, or philosophical training.<sup>13</sup> They also report that Seyedsayamdost (unpublished) tried to thoroughly replicate all the results presented by WNS with British participants, but he found no ethnically correlated differences. [Turri \(2013\)](#) also reports a lack of difference in evaluations of Gettier cases in tests with Western participants and participants from South Asia. Thus, a number of follow-up studies provide us with no evidence to support the phenomenon of cultural variation reported by WNS.<sup>14</sup>

Moreover, WNS only found a statistically significant difference in a single Gettier case on a single multicultural pool of American undergraduates, whereas they found no such difference between Westerners and East Asians in two additional experiments designed to elicit Gettier intuitions ([Weinberg et al. 2001](#), p. 30). The cases they investigated were not tested repeatedly, and they did not test other Gettier cases with different content. WNS also did not find a statistically significant difference in two of the three tested versions of the Truetemp case. Even Stich himself now questions the veracity of the data presented by WNS, calling it “a crude study” and admitting he’s “not sure [he] trust[s] it”.<sup>15</sup>

There are good reasons to doubt that the findings reported by WNS are robust; however, I do not want to rest my case on the current empirical data. Thus, I will suppose for the sake of argument that new empirical evidence might confirm that there is cultural variability. I will assume that this data has been impeccably produced, robustly replicated, and so we can confidently take it for granted that cultural variation exists.

<sup>13</sup> Moreover, [Nagel et al. \(2013, p. 4\)](#) points out that only two of the eleven reported comparisons by WNS found that the majority of one group went one way (i.e. “only believes”), and a majority of the other group went another way (i.e. “really knows”). In other words, actual group disagreement was reported only for two of the eleven comparisons discussed.

<sup>14</sup> In an unpublished study, Edouard Machery and his team report the first cross-linguistic evidence that Gettier intuitions are found across languages. This suggests that Gettier intuitions may be the expression of a core folk epistemology.

<sup>15</sup> Stich made these remarks during his talk at the Moral Sciences Club at the University of Cambridge on 14 May 2013.

If this were so, we might have a conceptual difference. For example, if the majority of Westerners refuse to apply the word ‘know’ to the luckless fall guy in a Gettier scenario, whereas the majority of South Asians are willing to attribute knowledge to him, then the concept expressed by the word ‘know’ arguably differs for these groups. Moreover, if another culture has a word that we have been translating as ‘know’ that lacks a similar extension, then perhaps their word should not be translated as ‘know’. Do these conceptual differences threaten the possibility of deriving epistemological insights from facts about the word ‘know’ or the concept it expresses? I’ll argue that they do not.<sup>16</sup>

Whether or not ‘knowledge’ is used to express several concepts will depend on one’s account of concept individuation. The issue of whether people who assign slightly different extensions to the word ‘knowledge’ have different concepts is vexing because “there is no consensus, either among philosophers or among linguists, about what constitutes a difference in meaning, or about how concepts are to be individuated” (Stich 2013, p. 5). According to Jackson (1998, 2005, 2011), if two people disagree about whether the victim of a Gettier scenario qualifies as knowing, then they simply have different concepts of knowledge and mean different things by ‘know’. If, however, one thinks that concepts are individuated in a less fine-grained way (e.g. Fodor 1998), then the universality of the concept expressed by the word ‘know’ is compatible with considerable culturally correlated variation.

The hypothesis I want to explore in the remainder of this paper is that there is a *universal core* to the concept of knowledge that is culturally invariant. According to this view, there is a core set of epistemic judgments (or “intuitions”) that do not vary cross-culturally. If this view is correct, then we have a strong reason to regard the concept of knowledge as universal. Moreover, if cross-cultural judgments about whether an agent knows sufficiently overlap, then epistemologists could still derive significant theoretical payoff by investigating the word ‘know’, the concept it expresses, and sentences of the form “S knows that *p*”. Differences in the extension of the word ‘know’ that exist at the periphery of our concept should not bar epistemologists from studying the deeper structural features of the concept of knowledge.<sup>17</sup>

<sup>16</sup> One response is to endorse conceptual pluralism. Sosa (2007, 2009) argues that cultural variation of epistemic judgments reveals that different cultures employ different standards of epistemic evaluation. Similarly, Jackson (1998, 2005, 2011) maintains that people who have different epistemic judgments about Gettier cases should be regarded as having different epistemic concepts. However, this type of epistemic pluralism is open to a serious challenge. If we agree that there are different concepts of knowledge, this leads us to wonder why we should care about the concept of knowledge that happens to be expressed by the English word ‘know’ (or the concept expressed by that word for some cultural subgroup). Without some reason to think that what one group calls ‘knowledge’ is any more valuable, desirable, or useful than what another group calls ‘knowledge’, it is hard to see why we should care if we can’t have it (Nichols et al. 2003, p. 245). To the extent that epistemologists remain interested in deriving theoretical insights by investigating the word ‘know’, any attempt to explain diverse epistemic judgments in terms of a diversity of concepts might undermine the common methodological approach. Thus, this type of pluralism about the concept of knowledge should make epistemologists deeply uncomfortable.

<sup>17</sup> Nagel (2012) has a similar view, but she defends this idea in a very different way.

## 5 The universal core of knowledge

The hypothesis that there is a universal core to the concept of knowledge deserves to be investigated seriously, especially since there is already evidence to support this view. For example, the study conducted by WNS found that for one crucial probe there was no statistically significant difference among *any* of the groups they examined. WNS tested the following ‘coin flip’ case:

*Coin Flip.* Dave likes to play a game with flipping a coin. He sometimes gets a “special feeling” that the next flip will come out heads. When he gets this “special feeling”, he is right about half the time, and wrong about half the time. Just before the next flip, Dave gets that “special feeling”, and the feeling leads him to believe that the coin will land heads. He flips the coin, and it does land heads. Did Dave really know that the coin was going to land heads, or did he only believe it? (Weinberg et al. 2001, p. 36)

There was widespread consensus amongst participants that Dave only believed (did not know) that the coin was going to land heads. If almost everybody agrees that we should not treat cases of subjective certainty as knowledge, then at least *some* of our epistemic judgments are universal. WNS are themselves open to the possibility of such a universal core to the concept of knowledge; they write, “the fact that subjects from all the groups we studied agreed in not classifying beliefs based on ‘special feelings’ as knowledge suggests that there may be a universal core to folk epistemology” (2001, p. 36).

Wright (2010) provides additional support for the hypothesis that there is a universal core to knowledge. Using a case developed by Swain et al. (2008), Wright gave the following scenario to a group of 188 undergraduate college students from the University of Wyoming:

*Chemist Case.* Karen is a distinguished professor of chemistry. This morning, she read an article in a leading scientific journal that mixing two common floor disinfectants, Cleano Plus and Washaway, will create a poisonous gas that is deadly to humans. In fact, the article is correct: mixing the two products does create a poisonous gas. At noon, Karen sees a janitor mixing Cleano Plus and Washaway and yells to him, “get away! Mixing those two products creates a poisonous gas!”

When asked whether Karen knew that mixing those two products would create a poisonous gas, the vast majority of participants said yes.<sup>18</sup> Swain, Alexander, and Weinberg observed the same results.

This data suggests that there may be a universal core to the concept of knowledge that is culturally invariant. Unfortunately, we do not currently have enough empirical data to demonstrate that the epistemic judgments of the vast majority of people will converge over a wide range of cases. This is largely because experimental epistemologists have been investigating a very limited range of examples—

<sup>18</sup> Wright found that her participants widely agreed on several cases (2010, pp. 501–502). However, the participants were predominately Caucasian, so this study is not ideal for my purposes.

Gettier scenarios, Truetemp cases, and skeptical pressure cases involving far-fetched possibilities that have not been eliminated by our evidence (e.g. that a certain zebra-looking animal is actually a cleverly painted mule).<sup>19</sup> It is not surprising that they have focused on these examples. Many major movements in epistemology over the past 50 years have relied on the judgments elicited by these thought experiments. Thus, experimental epistemologists are right to question whether the judgments allegedly held about these cases are in fact widely shared.

However, the cases tested by experimental epistemologists are an unrepresentative lot; they are not characteristic of the situations humans tend to encounter in daily life. People do not commonly need to determine whether the unlucky victim of a Gettier scenario knows, or whether someone who has miraculously had his brain reliably rewired is someone on whom we should rely, or whether far-fetched skeptical possibilities are actualized. Gettier scenarios, Truetemp cases, and skeptical possibilities are not the kinds of situations we routinely face in everyday life, so it would be surprising if our concept of knowledge was built to deal with them. Our practice of epistemic evaluation has not developed to regulate these types of situations. Thus, it is unsurprising that these conceptual differences—assuming they are genuine—have gone unnoticed, since the differences revealed by these cases are slight and the hypothetical cases that bring them out are unusual.

Swain et al. (2008) resist this line of reasoning. They argue that

as increasing amounts of empirical data raise trouble along different dimensions for intuitions about thought-experiments of various sorts, it becomes less reasonable to suppose that the intuitions experimental philosophers attack are an unrepresentative lot. (Swain et al. 2008, p. 148)

Their point is that as more empirical studies detect variation in our epistemic judgments, it becomes less reasonable to suppose that there are only a few circumstances under which our judgments display variation.

However, the empirical case for cultural variation has focused only on a few cases that, while popular in some corners of epistemology, are not the sorts of cases for which people commonly use the word ‘know’ in epistemic evaluation. For this reason, the existing data does nothing to show that cross-cultural divergence in epistemic judgments is a widespread phenomenon. The findings presented by WNS are compatible with convergence across a vast range of cases. It is plausible that the degree of cultural variability correlates with how paradigmatic the case in question is: the more a case represents a clear instance of the concept of knowledge, the less likely it is that there will be cultural divergence. Some cases strike us as clear examples of knowledge (e.g., beliefs gained through direct perceptual observation under ideal conditions), some cases strike us as clearly not knowledge (e.g., a lucky guess), and for other cases there may be no widely accepted answer (e.g., Gettier cases, Truetemp cases, and skeptical pressure cases) (Wright 2010, p. 500).

<sup>19</sup> Other examples in the literature display this limited focus. Nichols et al. (2003) and Buckwalter and Stich (forthcoming) test brain-in-a-vat cases, whereas Swain et al. (2008) test Truetemp cases.

As a result, the evidence for variation in our judgments about knowledge does not undermine the hypothesis that there is a universal core that is worth investigating. We can therefore deny the claim that “*basic* ideas about what is required for knowledge are apparently different across cultures” (Knobe and Nichols 2008, p. 11, emphasis mine).

If the concept of knowledge has a universal core, what does this core look like? The pragmatic theory outlined in Sect. 3 provides us with the answer: the core of knowledge is sufficiently reliable true belief, where ‘sufficiently reliable’ is unpacked roughly as ‘reliable enough to serve the purposes of members in the linguistic community’. An agent with a sufficiently reliable true belief is in a strong enough epistemic position with respect to some proposition  $p$  to eliminate all of the not- $p$  possibilities that are relevant to members of the community that might draw on her true belief. The relevant alternatives are those that are widely (but tacitly) agreed to be the ones that an informant must rule out. As I’ve argued, it is reasonable to presume that such judgments will sufficiently coincide to give us a set of core not- $p$  alternatives.<sup>20</sup>

## 6 Normative epistemology

Even if no actual society has a very different concept of knowledge, it is possible that some society could have had one. Imagine a society in which nothing remotely similar to our concept of knowledge is expressed by a word in their language; for example, it is conceivable that for some society the word best translated as ‘know’ is one that means ‘true belief’ or ‘justified false belief’. Consequently, you might worry that it is arbitrary to be interested in the concept we actually have, since it is possible that some society got along without it.

But not all epistemic concepts are created equal. We have the concept of knowledge to satisfy basic communicative and epistemic needs, such as the needs to identify, pool, and share reliable information. This explains why we should expect this concept to arise in any human society. Insofar as this concept has arisen to satisfy these basic pragmatic needs, it is plausible that any society lacking this concept would be epistemically impoverished. The concept of knowledge is deeply important to human life and thought; thus, any society that lacked this concept would find it difficult, if not impossible, to satisfy the same basic needs that the concept of knowledge is used to satisfy. Indeed, the very possibility of a well-developed human society getting by without a concept like knowledge is dubious if my account is correct.

Epistemologists can therefore go beyond the descriptive project, which reveals the present shape of our epistemic concepts, and engage in a normative project, which evaluates how well or poorly people’s epistemic practices satisfy their needs and goals. Putative facts about the roles of the concept of knowledge can be put to work for normative theorizing; we can criticize and reject the epistemic standards or concepts of a society that fail to satisfy basic human needs and goals. This would not be arbitrary,

---

<sup>20</sup> Elsewhere I have discussed the implications of this view for debates centering on epistemic contextualism (see Hannon (2013, Forthcoming)). Here I leave this matter open.

xenophobic, or unjustified. There are clear pragmatic grounds on which we can confidently judge our practices of epistemic evaluation as superior to a society that lacked a (near-) synonym for ‘knowledge’. Any such society would be unable to effectively meet our fundamental communicative and epistemic goals, and thus disadvantaged as a result.<sup>21</sup> Moreover, this is an area of philosophical inquiry where experimental philosophy has no bearing. The results of recent experimental studies only reveal how we in fact categorize cases of knowledge, not how it would be good to categorize them (i.e. how we ought to categorize them for our purposes or benefit).

If our epistemic judgments about Gettier cases, Truetemp cases, and skeptical pressure cases were culturally variable, this would show that the universal core of knowledge does not include a number of epistemic judgments that have played a central role epistemology. Is this a problem for epistemologists who want to rely on these judgments when theorizing? Yes, it certainly is. Let’s consider Gettier cases for illustrative purposes. If there were genuine cultural variability about whether the victim of a Gettier scenario knows, then epistemologists who seek to solve the Gettier problem would really only be investigating a concept of knowledge that is possessed by some people (i.e. Westerners) and not others (i.e. East Asians). The Gettier problem would simply not arise for those cultures or dialect groups that do not have the concept possessed by the majority of Western participants in WNS’s study. The same could be said for Truetemp cases and skeptical pressure cases. (I have argued in Sect. 4.2 that there is no robust data in favor of such cultural variability, but I am granting the existence of such data for the sake of argument.)

However, the fact that some limited range of epistemic judgments are shown to be culturally variable does not threaten my epistemological project, which aims to show that there is a universal core to the concept of knowledge. Moreover, anyone engaging in this type of project can still derive important theoretical insights by investigating uses of the English word ‘know’. The empirical data are only worrying if we are committed to research programs that mistakenly assume everybody will agree that the protagonists of Gettier cases, Truetemp cases, and skeptical pressure cases lack knowledge. But the instability of some epistemic judgments under some circumstances is not a good reason to reject epistemic judgments as evidence more generally. *Contra Stich (1988)*, the mere possibility of cultural diversity is not enough to undermine analytic epistemology.

This raises an important question: which are the philosophically interesting cases? You might think that philosophy is most interesting (and of most value) when it is working “at the margins”, wrestling with unclear and borderline cases (Wright 2010, p. 500). If this were true, then cultural variation could undermine a potentially important area of philosophical inquiry. However, my proposal changes the nature of the debate between epistemologists who study facts about the word ‘know’ and their critics. One goal of epistemology is (or should be) to elucidate the deeper core of knowledge. The core presents features of the concept that are necessary if it is to satisfy

---

<sup>21</sup> There may be room for improvement. I do not insist that our concept of knowledge is the best one to have all things considered. My claim is that the concept of knowledge is good to have because it does important work and solves coordination problems centering on the basic needs to identify, pool, and share reliable information.

certain pragmatic needs, whereas the periphery has features that are not necessary to satisfy such needs. One of the virtues of my approach is that it separates those features of a concept that emerge from the purpose of its origin from those features that are a contingent and peripheral matter of history. According to my proposal, the common methodological view remains sound and suitable for elucidating the core of the concept of knowledge.<sup>22</sup>

Do we know which epistemic judgments display variation and which do not? Even if our epistemic judgments can in principle be used as evidence in philosophy, the current empirical data might be taken to suggest that, at this time, we cannot tell which judgments are universal and which are variable. Swain, Alexander, and Weinberg go so far as to say that “there is unlikely to be a fixed set of intuitions about a particular thought-experiment to which we can appeal” (2008, p. 141).

However, Wright (2010, 2013) has shown that people are significantly more confident in their knowledge-ascribing and knowledge-denying judgments for cases where agreement is widespread. This suggests that people have introspective access to certain attitudinal states that can serve as a reliable indicator of our core epistemic judgments (Wright 2013, p. 8). High degrees of confidence track widespread agreement about whether a subject knows, and low degrees of confidence track cases where disagreement is more likely. Indeed, this is exactly what we should expect if my account is correct. People will have the competency needed to distinguish those judgments that are more central to our concept of knowledge from those that are not.

## 7 Conclusion

Reflecting on why humans think and talk about knowledge provides us with independent contrary evidence to the view that the concept of knowledge is culturally variable. This shields much epistemology from the recently popular objection that we cannot derive important theoretical insights by investigating the English word ‘know’, the concept it expresses, and/or the semantic properties of knowledge sentences. If the concept expressed by the English word ‘know’ is universal, then epistemologists can investigate a concept that we all share despite our diverse cultural and linguistic backgrounds. In this way, my proposal vindicates a common methodology in epistemology.

Moreover, we need not deny that there is (or could be) cross-cultural variation in epistemic judgments. Variation in some areas is compatible with uniformity in others. This leaves space to explore systematic variation in the extension of the word ‘know’ as used by English speakers from different cultural backgrounds, as well as the ways in which the word best translated as ‘know’ in other languages has a different extension from the English word.

---

<sup>22</sup> Precisely why people from different cultural backgrounds would assign different extensions to the concept of knowledge is not something I will attempt to answer here (an explanatory hypothesis is not needed for my purposes). The method of practical explication leaves gaps that are filled in different ways, in different societies, and at different times. We can speak of a core concept of knowledge that, in different times and places, receives different forms of cultural elaboration.



**Acknowledgments** Thanks to Mikkel Gerken, Stephen Grimm, and Robin McKenna for comments on an earlier draft.

## References

- Amberber, M. (Ed.). (2007). *The language of memory in a cross-linguistic perspective*. Berlin: Mouton de Gruyter.
- Austin, J. L. (1961). *Philosophical papers*. Oxford: Clarendon Press.
- Beebe, J. (2012). Social functions of knowledge attributions. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.
- Boyd, K., & Nagel, J. (2014). The reliability of epistemic intuitions. In E. Machery & E. O'Neill (Eds.), *Current controversies in experimental philosophy*. London: Routledge.
- Buckwalter, W., & Stich, S. (2011). Gender and the philosophy club. *The Philosopher's Magazine*, 52, 60–65.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Oxford University Press.
- Dancy, J. (1992). Review of knowledge and the state of nature. *The Philosophical Quarterly*, 42(186), 393–395.
- DeRose, K. (1995). Solving the sceptical problem. *Philosophical Review*, 104(1), 1–52.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67(24), 1007–1023.
- Fodor, J. (1998). *Concepts: Where cognitive science went wrong*. Oxford: Oxford University Press.
- Gerken, M. (2013). The roles of knowledge ascriptions in epistemic assessment. *European Journal of Philosophy*, 21(2).
- Goddard, C. (2010). Universals and variation in the lexicon of mental state concepts. In B. Malt & P. Wolff (Eds.), *Words and the mind*. Oxford: Oxford University Press.
- Goddard, C., & Wierzbicka, A. (Eds.). (2002). *Meaning and universal grammar—Theory and empirical findings* (Vol. I and II). Amsterdam/Philadelphia: John Benjamins.
- Hannon, M. (2013). The practical origins of epistemic contextualism. *Erkenntnis* 78(4), 889–919.
- Hannon, M. (2014). Fallibilism and the value of knowledge. *Synthese* 191(6), 1119–1146.
- Hannon, M. (Forthcoming). Stabilizing knowledge. *Pacific Philosophical Quarterly*.
- Haspelmath, M., & Tadmor, U. (2009). World loanword database. Max Planck Digital Library. <http://wold.livingsources.org/>.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Jackson, F. (2011). On Gettier holdouts. *Mind and Language*, 26(4), 468–481.
- Jackson, F. (2005). Ramsey sentences and avoiding the sui generis. In H. Lillehammer & D. H. Mellor (Eds.), *Ramsey's legacy*. Oxford: Clarendon Press.
- Jackson, F. (1998). *From metaphysics to ethics: A defense of conceptual analysis*. Oxford: Clarendon Press.
- Lackey, J. (2012). Group knowledge. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.
- Lehrer, K. (1990). *Theory of knowledge*. Boulder: Westview Press.
- Lycan, W. (2006). On the Gettier problem problem. In S. Hetherington (Ed.), *Epistemology futures*. Oxford: Oxford University Press.
- Machery, E., Stich, R., Chatterjee, A., Karasawa, K., Struchiner, N., Usui, N., & Hashimoto, T. *Gettier across cultures*. Unpublished Manuscript.
- Nagel, J. (2012). Intuitions and experiments: A defense of the case method in epistemology. *Philosophy and Phenomenological Research*, 85(3), 495–527.
- Nagel, J., Juan, S., & Mar, R. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129(3), 652–661.
- Nichols, S., Stich, S., & Weinberg, J. (2003). Metaskepticism: Meditations in ethnoepistemology. In S. Luper (Ed.), *The skeptics*. London: Ashgate.
- Russell, J. A. (1991). Culture and the categorization of emotion. *Psychological Bulletin*, 110, 426–450.
- Rysiew, P. (2012). Epistemic scorekeeping. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- Seyedsayamdost, H. *On normativity and epistemic intuitions: Failure to detect differences between ethic groups*. Unpublished Manuscript.

- Sosa, E. (2009). A defense of the use of intuitions in philosophy. In Dominic Murphy & Michael A. Bishop (Eds.), *Stich and his critics*. Chichester: Wiley-Blackwell.
- Sosa, E. (2007). Experimental philosophy and philosophical intuition. *Philosophical Studies*, 132(1), 99–107.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124(3), 272–283.
- Stich, S. (2013). Do different groups have different epistemic intuitions? A reply to Jennifer Nagel. *Philosophy and Phenomenological Research*, 87(1), 151–178.
- Stich, S. (1988). Reflective equilibrium, analytic epistemology, and the problem of cognitive diversity. *Synthese*, 74(3), 391–413.
- Swain, S., Alexander, J., & Weinberg, J. (2008). The instability of philosophical intuitions: Running hot and cold on Truetemp. *Philosophy and Phenomenological Research*, 76(1), 138–155.
- Turri, J. (2013). A conspicuous art: Putting gettier to the test. *Philosophers' Imprint*, 13(10), 1–16.
- Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy* (pp. 17–45). Oxford: Oxford University Press. (Reprinted from *Philosophical Topics*, 29(1–2), 429–460, 2008).
- Wierzbicka, A. (1999). *Emotions across languages and cultures*. Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, J. (2013). Tracking instability in our philosophical judgments: Is it intuitive? *Philosophical Psychology*, 1–17.
- Wright, J. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, 115(3), 491–503.