

A defence of conceptual analysis as a linguistic endeavour

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Grindrod, J. ORCID: <https://orcid.org/0000-0001-8684-974X>
(2023) A defence of conceptual analysis as a linguistic endeavour. *Theoria*, 89 (4). pp. 516-534. ISSN 1755-2567 doi: <https://doi.org/10.1111/theo.12477> Available at <https://centaur.reading.ac.uk/111927/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/theo.12477>

Publisher: Wiley-Blackwell

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

A defence of conceptual analysis as a linguistic endeavour

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Grindrod, J. ORCID: <https://orcid.org/0000-0001-8684-974X>
(2023) A defence of conceptual analysis as a linguistic endeavour. *Theoria*. ISSN 1755-2567 doi:
<https://doi.org/10.1111/theo.12477> Available at
<https://centaur.reading.ac.uk/112285/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/theo.12477>

Publisher: Wiley

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

ORIGINAL ARTICLE

A defence of conceptual analysis as a linguistic endeavour

Jumbly Grindrod 

Philosophy Department, University of Reading, Reading, UK

Correspondence

Jumbly Grindrod, Philosophy Department, Edith Morley, Whiteknights, University of Reading, Reading, RG6 6AH, UK.
Email: j.grindrod@reading.ac.uk

Abstract

In this paper, I outline and defend a traditional yet controversial view of conceptual analysis, particularly as it is used in epistemology. I will defend the view against a number of objections, all of which focus on the idea that conceptual analysis relies upon *linguistic* intuitions. Rather than trying to deny this claim, I will seek to vindicate the use of conceptual analysis within epistemology even given its reliance on linguistic intuitions. To do so, I first outline the view of conceptual analysis I have in mind and then respond to three broad objections from Quill Kukla, Allan Hazlett and Stephen Stich.

KEYWORDS

conceptual analysis, epistemic normativity, experimental philosophy, knowledge, linguistic intuitions

1 | INTRODUCTION

In ‘Postwar Oxford philosophy’ (Grice, 1989), Grice outlines and defends the project of conceptual analysis and in particular the extent to which it appeals to ordinary language use. Running throughout the paper is a clear desire to respond to the worry that conceptual analysis is merely a linguistic task and not a philosophical one. Some of the objections Grice considered in writing that paper are still being considered today, albeit in slightly different form. I agree with Grice’s own verdict of the paper that his ‘nose was pointed in the right direction’ (Grice, 1989, p. 181), and, like Grice, I intend to outline and defend conceptual analysis as a philosophical method. Specifically, I will be focused on conceptual analysis as it is used in epistemology. Also like Grice, I am concerned with the objection (roughly) that conceptual analysis is a linguistic methodology that is unsuitable to the aims of philosophical inquiry. This concern has been developed in a few different ways across the years, and in this paper, I will focus on arguments put forward by Kukla (2015), Hazlett (2018) and Stich (1988, 2009). There are points of overlap across the three authors, but also important points of difference, and so we will benefit from

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Author. *Theoria* published by John Wiley & Sons Ltd on behalf of Stiftelsen Theoria.

considering the arguments of each author at a time in order to present a clear defence of conceptual analysis as a methodological tool within epistemology.

The picture that will emerge will be self-consciously linguistic insofar as the intuitions relied upon within conceptual analysis are, I argue, rightly understood as linguistic intuitions. In that respect, we must acknowledge that the outcome of any given conceptual analysis should (initially at least) be restricted to a given language. Adopting this linguistic picture of conceptual analysis does have ramifications on how we are to understand the relationship between epistemic evaluative terms (such as ‘knowledge’ and ‘justification’) and epistemic normativity more widely. Crucially, conceptual analysis is an important project (or set of projects) within epistemology, but epistemology is not restricted entirely to conceptual analysis. In order to develop these remarks further, I will first outline the picture of conceptual analysis defended in this paper, and then consider the objections against such a picture from Kukla (in Section 3), Hazlett (in Section 4), and Stich (in Section 5).

2 | CONCEPTUAL ANALYSIS AS TARGETING LEXICALISED CONCEPTS

In employing conceptual analysis, we take a notion of philosophical interest and attempt to provide some account of its application conditions, that is, when it does and does not apply. Traditionally, such conditions were viewed in terms of necessary and sufficient conditions, but that picture is not crucial to the view defended in this paper. It may be instead, that we provide only necessary or sufficient conditions, or that we capture the notion in terms of a prototype, or find some other way of explaining why the notion applies in certain instances, but not in others.¹ In embarking on this project, we rely upon our intuitions regarding the target notion and whether it is present or absent in the hypothetical scenario we are considering.

To say this much is to say very little about what type of thing a ‘notion’ actually is. Let us follow Goldman (2010) in instead using the general term ‘classifier’. Classifiers are:

...things including properties, universals, kinds (natural and otherwise), meanings, individual concepts, that either have instances and counter-instances, or are standards or criteria that fix or pick out or determine sets of instances or counter-instances. (Goldman, 2010, p. 125)

Goldman distinguishes between *person-specific classifiers*, *community-specific classifiers* and *free-floating classifiers*. Person-specific classifiers are the classifiers associated with a given individual. For instance, you may be interested in investigating the concept that I have stored under the lexical entry of ‘modern art’, what I intend to express when I use the term, and where this concept may be different to other people’s. Alternatively, you may be interested in investigating community-specific classifiers: the concepts or linguistic meanings that are adopted by a given community. Finally, you may be interested in investigating the classifiers that are tied to no particular individual or set of individuals, for instance, if you are investigating universals or Platonic forms.

Given this three-way distinction, we can then ask: which of these three kinds of classifier should serve as the topic of inquiry in philosophical (and specifically epistemological) conceptual analysis? In asking this question, we need to consider not only what kind of classifier is of philosophical interest, but whether it is feasible to investigate that kind of classifier via the method of conceptual analysis. While free-floating classifiers, should they exist, certainly seem worthy of philosophical attention, it seems hard to justify the use of conceptual analysis in

¹Thanks to an anonymous referee for emphasising this point.

investigating them. As Goldman (2010) argues, the idea that our intuitions regarding particular cases should serve as evidence for the nature of free-floating classifiers looks implausible without some story of why it is that our intuitions do track their nature in the way required. On the other hand, person-specific classifiers seem suitable for investigation via conceptual analysis but arguably seem less philosophically interesting. If I treat the epistemological project of understanding knowledge as merely the project of understanding what concept I have stored under the lexical entry for ‘knowledge’—whether or not it matches up well with the rest of my community—the resulting project seems far-removed from the project usually engaged in within epistemology in both ambition and method.²

Community-specific classifiers fare better than person-specific classifiers in this regard. For instance, in considering the concept of knowledge that is prevalent within a community, we are no longer considering the potentially idiosyncratic concept of a given individual. In that respect, the project has the more universal flavour that philosophers seek. At the same time, it also seems reasonable to consider our intuitions regarding the correct application of these concepts as good defeasible evidence of what their correct application actually is. In virtue of being members of our community and competent speakers of the language, it seems *prima facie* plausible that we have a good grasp on the classifiers at play within our community, even if we are on occasion prone to error. So community-specific classifiers seem to be the most plausible classifier-type for the target of philosophical investigation via conceptual analysis.

The central case of community-specific classifiers are linguistic meanings. We take the community we are interested in to be a linguistic community, and we investigate the meanings associated with an expression in the community’s language. Given that members of the community are competent speakers of the language, it is reasonable to suppose that their intuitions regarding whether a given expression applies or not in a particular situation is good defeasible evidence for whether the expression does in fact apply or not. There may be community-specific classifiers that are not associated with a term in the language, and perhaps instances where new expressions are coined are often cases where a community-specific classifier finally gets named, thus giving members of the community the opportunity to easily talk about what falls under that classifier without resorting to descriptive or demonstrative terminology. Still, the typical case will be lexicalised concepts, and this does seem to be the case in many philosophical projects, where we consider the extension of notions like ‘good’, ‘know’, ‘believe’ and so forth.

Conceptual analysis understood in this way has been subject to a number of different attacks in recent years. I will briefly mention three broad forms of objection, although what I have to say in response to them is not particularly original. First, some (Cappelen, 2012; Deutsch, 2010; Williamson, 2007) have argued that it is a mistake to think that philosophy relies on intuitions in the manner described above. Beyond the fact that the so-called *intuition-denying* position provides a way of avoiding the threat posed by results in experimental philosophy (discussed below), a few different arguments for the view have been given. They include:

- i. That if intuitions were used as evidence, then there would be a problematic gap in philosophical argumentation between the psychological premises about our intuitions and the non-psychological conclusions about some philosophical topic (Williamson, 2007).
- ii. That the idea that intuitions do play a role is motivated by a misguided appeal to a principle of evidence neutrality: roughly, that all party to a debate can in principle agree on what the evidence is, even if they cannot agree on the conclusions to be drawn from it (Williamson, 2007).
- iii. That inspection of previous philosophical works reveals little or no appeal to intuition (Cappelen, 2012; Deutsch, 2010).

²Grice discusses this kind of project in considering the possibility that his application of a concept could be idiosyncratic, and is keen to emphasise that there would still be benefit to it: ‘even if my assumption that what goes for others is mistaken, it does not matter; my philosophical puzzles have arisen in connection with my use of E, and my conceptual analysis will be of value to me (and to any others who may find that their use of E coincides with mine)’ (Grice, 1989, p. 175).

As stated, it is not the aim of this paper to respond to the intuition-denying view, but I will give some idea as to how a response might go. Regarding (i) and (ii), both arguments have less force if we take conceptual analysis to be targeting our lexicalised concepts. On (i), we can resist the idea that our conclusions in conceptual analysis are non-psychological. Instead, the desired conclusions concern the application conditions of our lexicalised concepts, which is partly a psychological matter.³ Argument (ii) is simply resisted by denying that the role of intuition in conceptual analysis is motivated (even implicitly) by the evidence neutrality principle. As suggested above, the role of intuition is not motivated by theoretical disagreements, but by the plausible enough idea that there is a reliable connection between the application conditions of the lexicalised concept and the application of that concept to hypothetical cases. Regarding (iii), I basically disagree with the interpretation of the previous literature. I find it hard to understand the nature of the Gettier debate, for instance, without identifying some role for intuition in considering the thought experiments. Deutsch and Cappelen are right to point out that oftentimes the arguments used in such debates are more sophisticated than just a bare appeal to intuitions. But, like Chudnoff (2017), I am highly sceptical that this strategy could plausibly exorcise intuitions from the debate altogether.⁴

Secondly, some have argued that while conceptual analysis may depend on intuitions, those intuitions are not linguistic in the manner I have suggested. For instance, Sosa (2009) explicitly resists the idea that in considering the nature of knowledge, we are thereby considering the meaning of ‘knowledge’:

The question is not just whether ‘knowledge’ applies to the protagonist in a certain example. The question is whether the protagonist who satisfied the conditions specified in the example would *know*. To see that this is the interesting question in epistemology, we need only retreat to our own reflection, leaving behind any kind of dialogue, whether in journal, conference, seminar, or hallway, and just entertain the question reflectively *in foro interno* [...] We can of course consider *also* whether in our idiolect of the moment it would be correct to apply our word ‘knows’ to such a justified believer of a truth. But this is a different question, though one now with an equivalent answer. (Sosa, 2009, p. 104)

Sosa is obviously correct that there are two distinct questions to be asked here, one about knowledge and one about ‘knowledge’. But the important question is whether, in introspecting on the correct answer to each of these questions, we are drawing on the same set or distinct sets of intuitions. I am inclined to think we are drawing upon the same set of intuitions. It might be objected that when we consider the application conditions of a term, we must be sensitive to further considerations, for example pragmatic considerations regarding what a speaker would communicate were they to use that term within a given context. In that respect, the objection goes, the two questions that Sosa identifies are differentiated in terms of the intuitions used to

³As Brown (2011, p. 505, fn. 15) notes in passing, this is essentially the line taken by Goldman and Pust (1998).

⁴For further criticisms of the intuition-denying position, see Brown (2011) and Nado (2016). Williamson has said more on this topic than I can do justice to here, particularly in his work of 2007. I will mention, however, that it is often the case that Williamson is considering questions only indirectly related to the position defended here. For instance, sometimes he considers whether *all* philosophical questions could be about linguistic or conceptual issues (2007, ch. 2), whereas I am only concerned with conceptual analysis as one type of philosophical methodology. Elsewhere, he considers whether the judgements used in philosophical thought experiments are different in kind to our everyday judgements (2007, ch. 6), arguing that judgements in thought experiments employ the same counterfactual imagining abilities that we use in ordinary life. I agree with Williamson about the broader point here, that the intuitions garnered from thought experiments are not peculiar to philosophical investigation. Like Grice, I do not take the kind of conceptual analysis procedure defended here to be restricted to philosophical matters. We often do have to consider the meaning of a term in everyday life, and in doing so we resort (even if roughly) to an inspection of the concepts we are employing.

answer them. In fact, however, this is no difference at all. For pragmatic considerations must also be taken into account when theorising about knowledge in precisely the manner that Sosa has in mind. If in thinking through a hypothetical case we reach a verdict that, say, a subject does not know a certain proposition, we have to be wary of the possibility that our intuitions are really picking up on the fact that it would be pragmatically inappropriate to say that the person knows within that context. In this respect, there is no clear difference between our intuitions regarding the application of 'know' and our intuitions regarding whether someone knows.⁵ Of course, if Sosa or anyone can provide a way of systematically distinguishing between these questions in terms of the intuitions we draw upon, then this is good news for the health of conceptual analysis. But I am sceptical that we can be so discerning in selecting our intuitions for use in conceptual analysis.

Thirdly, the picture of conceptual analysis given above has been challenged by the so-called *negative research programme* in experimental philosophy, which aims to raise doubts regarding conceptual analysis as an armchair method of theorising about philosophical topics. The clearest way to do this is to show that the intuitions garnered by philosophers are unreliable insofar as they can be affected by non-relevant factors such as the demographic, ethnicity or language of the intuiter, ordering effects, framing effects, font used and so forth. The force and success of the negative research programme deserves greater attention than I will give it here (see Weinberg, 2017), but I will just note two points that I think should prevent one from thinking that what occurs in the remainder of this paper is uninteresting because the negative research programme has already shown conceptual analysis to be a busted flush. First, given that it is plausible to think that a given individual's intuitions regarding the applications of their own concepts is good (defeasible) evidence for the application conditions for that concept, the onus is very much on the negative research programme to show that such intuitions are problematic. With this in mind, Weinberg (2017) is one negative research programmer who has been careful to stress that they are not claiming that philosophical intuitions are problematic across the board. Instead, the findings to the effect that certain intuitions are sensitive to irrelevant factors should be treated, initially at least, as relevant only to the specific domain of inquiry tested. If the negative research programme continues to produce such results time and again, then we may reach a stage where the inductive leap would be appropriate to the claim that all such intuitions garnered via conceptual analysis are unreliable. While lots of interesting results do continue to be produced, I think it is fair to say that we are not at that stage yet. Turning to the domain of epistemology in particular, results from the negative research programme have been a little mixed. On the one hand, results like Weinberg et al. (2001), Swain et al. (2008) and Beebe and Buckwalter (2010) have shown that epistemic intuitions appear to be sensitive to ethnic background and socio-economic status, ordering effects, and moral considerations, respectively. On the other hand, some of these results have failed to be replicated (Kim & Yuan, 2015; Seyedsayamdost, 2015), while certain key epistemological intuitions (i.e., Gettier intuitions) have been shown to arise across a range of different languages (Machery et al., 2015). With this in mind, it is not yet plausible to accept the claim that epistemic intuitions of the kind usually garnered in epistemological conceptual analysis are unreliable. The evidence is not there yet.

⁵To take another case, Horne and Cimpian (2018) have presented experimental results suggesting that there are different epistemic standards associated with 'know' and 'knowledge', respectively, with the latter typically associated with a more demanding standard. Perhaps these findings could be interpreted as showing that the linguistic practices associated with 'know' should be separated from whether someone counts as possessing knowledge or not. But the idea that the use of 'knowledge' is a more reliable guide to the object level issue of what we know is obviously unmotivated. Rather, the natural way to interpret these findings (as Horne and Cimpian do) is as driving a wedge between what is picked out using the verb phrase compared to what is picked out using the noun phrase. These findings certainly do warrant further investigation, and a number of possible explanations await. For instance, there may be a general pragmatic effect associated with noun phrases rather than verb phrases (or vice versa). Thanks to an anonymous referee for raising these findings.

Of course, this issue in some sense presupposes that we can identify extraneous factors, but what independent basis do we have to rule out some factors as extraneous and rule in others as revealing something important about, for example, the knowledge concept? As Cummins (1998) stresses, in scientific practice the noise in observational data is identified (and, where possible, reduced) by calibrating the instrument with which we observe. Check how good your telescope is by checking how your telescope represents the properties of an object that you are already familiar with. But with intuitions, we seem to have no such calibration procedure available to us. However, rather than thinking of this as a knockdown objection to the use of intuitions in philosophy, I instead take this to be one of the difficulties that a theorist must take careful consideration of (no-one said this was going to be easy). A recent debate regarding the relevance of moral considerations in theorising about knowledge illustrates the point. As mentioned, Beebe and Buckwalter (2010) found that the moral valence of a side effect of some actions affect whether participants are willing to claim the actor knew the side effect would occur—a so-called *moral side-effect effect*. One way to interpret this result is as showing that knowledge intuitions are affected by moral valence as an extraneous variable. An alternative is to claim that whether we know depends on certain moral features of the situation (the *moral encroachment* position, which has found increasing support in recent years; Basu, 2019; Fritz, 2017). So how do we decide between the two approaches? A number of considerations come into play, including how disruptive the moral encroachment view would be in capturing other cases, but also whether there are general effects of moral features on our thinking. This is analogous to a theoretical issue that many philosophers of language are already acquainted with—considering whether some pattern in our linguistic intuitions should be captured in our semantic or pragmatic theory. In deciding this, the theorist will consider whether accommodating such intuitions does damage elsewhere to one’s semantic theory, but also whether there are known general pragmatic phenomena that may play some explanatory role.⁶

This paper will defend a view of conceptual analysis as a linguistic endeavour wherein we attempt to capture a lexicalised concept in circulation within a linguistic community. In the next section, I will first focus on Kukla’s objection that epistemic discourse is in fact too messy to admit of conceptual analysis.

3 | KUKLA ON THE MESSINESS OF EPISTEMIC DISCOURSE

The first objection to be addressed is that epistemic discourse is in fact too messy, varied and possibly even conflicting to be something that could serve as the basis of a conceptual analysis. This form of objection is not new—in fact it is one that Grice considered some form of (in his opponent’s voice):

Ordinary language suffers from various defects which unfit it for conceptual analysis, or at least prevent conceptual analysis, or at least prevent conceptual analysis from achieving any results which are worth the effort involved in reaching them. Such defects are ambiguity, misleadingness, vagueness, and the incorporation of mistakes or absurd assumptions. (1989, pp. 176–177)

Stated in this way, it is hard to see what the threat to the project is. As Grice notes, phenomena such as ambiguity and vagueness (and we may add context-sensitivity to this) are not barriers to providing a conceptual analysis, provided that the analysis can capture or at least accommodate the phenomena. Still, there may come a point where we would have to allow that a discourse is too messy to admit of any substantive analysis. Quill Kukla (2015, p. 205) has argued

⁶A similar point holds when considering the potential effect of heuristics and biases on our linguistic intuitions.

that 'in ordinary English, we use 'know' and its cognates idiomatically in a wide and conflicting variety of ways, many of which simply haven't been of interest to analytic epistemologists'. For instance, analytic epistemology, focusing largely as it has on propositional knowledge, will not easily be able to capture the following idiomatic uses of 'know' (Kukla, 2015, p. 206):

- iv. Do you know karate?
- v. I realised once we broke up that I never really knew her.
- vi. I don't know what you are talking about.

- vii. Your body knows what it needs.

Kukla suggests that analytic epistemology has employed an overly-narrow focus on know-that constructions and so largely ignores uses of 'know' such as (iv–vii). Once we take such uses into account, it no longer becomes clear that we would be able to provide an analysis of knowledge that captures even widely held assumptions within epistemology, for example that knowledge is a propositional attitude that bears some relation to belief.

Contrary to Kukla, however, I do not take idiomatic uses to be a cause for concern. Idiomatic uses are widespread across language, and while they certainly do complicate matters when it comes to investigating a particular part of language, it is not an insurmountable task to set them to one side. As an analogy, consider discourse about heat. As well as central cases—for example 'the pan was very hot', 'it is cold in here'—we get a range of idiomatic uses involving hyperbole, metaphor and other effects: 'the meeting was really heating up', 'it's freezing cold in here' and so forth. If I wanted to provide a theory about the way that heat is spoken about, or capture the application conditions of 'hot', 'cold' and so forth, I should not be put off by the fact that there are these idiomatic uses.

Furthermore, of the particular examples that Kukla appeals to (iv–vii) it seems that standard distinctions drawn in analytic epistemology can be used to make sense of them. The distinction between propositional knowledge, know-how and objectual knowledge is a well-worn one, and it is a distinction that helps make sense of (iv–vi). (iv) looks like a clear case of know-how, (v) looks like some form of objectual knowledge, and (vii)—a know-wh construction—can plausibly be captured in terms of propositional knowledge.⁷

So the response here is not that the messiness of epistemic discourse could not possibly render conceptual analysis unviable. It is that there is no reason to think that the messiness we find in ordinary epistemic discourse will prove to be a stumbling block to conceptual analysis.⁸

In a related point, Kukla states:

There are also uses of 'knowledge' that belong not to everyday vernacular language but to disciplines other than philosophy. Sociologists and anthropologists, for instance, will often refer to multiple 'knowledges'—subaltern knowledges, indigenous knowledges, and the like—meaning something like systems of inquiry, lore,

⁷Some have argued that know-wh is best understood as a tertiary knowledge relation that holds between a subject, a proposition and a question (Schaffer, 2007). But clearly this kind of account of know-wh still seems continuous with the general project of capturing the nature of propositional knowledge.

⁸In his earlier (2010) paper, Hazlett also points to a number of different uses that knowledge attributions are put to, and these could be added to Kukla's list in order to strengthen the case for messiness. However, like Hazlett (2010), I do not take the cases that he points to there to suggest that an analysis of knowledge is not possible – Hazlett instead uses those cases to argue for a non-factive account of knowledge. We could also add discourse-marking uses of 'know', e.g. 'You know, I never liked that movie'. However, as Hansen et al. (2021, p. 255) note in their corpus analysis of 'knows', 'The fact that [know] appears to have different meanings in its objectual, propositional, and discourse marking forms means that it is a relatively straightforward exercise to exclude the non-propositional meanings of [know] when investigating whether the meaning associated with the propositional use of [know] corresponds with the concept of knowledge of interest to epistemologists. The prevalence of discourse marking uses of [know] in non-philosophical talk therefore doesn't threaten ordinary language approaches in epistemology.' Thank you to an anonymous reviewer for emphasising this point.

and folk beliefs that are local in some way. Such a phrase may make philosophers cringe, but precisely the point here is that we cannot presuppose in advance that technical epistemological concepts of knowledge are the right ones, or that they map in any predictable way onto other sorts of discourse. (2015, p. 206)

The thought here is that epistemologists have been close-minded in thinking that the kind of knowledge they have typically been interested in is the only form of knowledge going, that in fact, when we look to knowledge discourse in other disciplines, we find a range of notions being referred to by the term ‘knowledge’. However, contrary to Kukla, there is good reason to think that philosophical conceptual analysis is tracking the ordinary notion of knowledge in a way that is not challenged by the fact that the term ‘knowledge’ may be used in other disciplines. In order for the distinctive use of the term in other disciplines to be of relevance to philosophy, it first has to be considered whether those disciplines, be they sociology, anthropology or whatever else, are using ‘knowledge’ as a theoretical term. The meanings of theoretical terms are stipulated according to the interests of the theorist, and so merely pointing to the fact that theorists in particular disciplines use ‘knowledge’ in a distinctive way provides no challenge to the conceptual analysis project of investigating the community-specific classifier. Furthermore, even if we allow that, for example subaltern knowledge, indigenous knowledge and so forth, are bona fide epistemic categories that have to be captured via our best theories of knowledge and justification, it remains to be seen that these categories cannot be captured via fairly standard epistemological theory. For instance, if we take subaltern knowledge to just be a set of propositions typically known only by the subaltern community, this is a notion that is easily captured via the analytic epistemologist’s standard toolkit. Simply put, it may be that such categories reduce into the categories typically discussed in analytic epistemology.⁹ So again, much more has to be done to show that epistemological conceptual analysis is rendered problematic by the above considerations.

4 | HAZLETT’S DILEMMA

Like Kukla, Hazlett (2018) argues that the conceptual analysis project in epistemology is problematic and also like Kukla, recommends that we instead adopt a more stipulative approach whereby we stipulate knowledge as a (or possibly the) valuable cognitive state and then investigate what such a cognitive state must look like in order for it to be valuable. However, as with Kukla, I will not focus on the prospects of the recommended stipulative approach and focus instead on whether Hazlett’s criticism succeeds. Hazlett argues that once we acknowledge that the intuitions we rely upon in conceptual analysis are linguistic, we face a dilemma. The dilemma arises when we ask whether in appealing to linguistic intuitions we should restrict ourselves to a single language or whether we should consider linguistic intuitions from across the full range of languages. He labels the former approach ‘parochialism’, and the latter approach ‘cosmopolitanism’, and he argues that either approach is problematic. I will argue that neither approach is problematic. I will turn to the cosmopolitan horn of the dilemma first.

According to the cosmopolitan approach, in attempting to capture application conditions for knowledge, we ought to take into account how speakers from a range of languages react to particular thought experiments. Hazlett presents a Quinian problem regarding translation that the cosmopolitan approach faces. Suppose we are investigating the predicate ‘swonk’ in some distinct language that is taken by the cosmopolitan theorist to translate to ‘know’. As a result,

⁹Note that further epistemological issues may arise when we focus, for example, on the subaltern community. For example, if there are a set of true propositions typically only known by a given community, this raises the partly epistemological question of why that is the case. But this new question does not show the prior question of the nature of knowledge to be flawed in any sense.

it falls within the cosmopolitan project to investigate whether, for example, 'swonk' is factive in the way that 'know' is commonly taken to be. Suppose that upon investigation, we conclude that 'swonk' is not factive:

viii. Sometimes, 'S swonk p' is true and $\neg p$.

The problem that Hazlett identifies is that rather than taking (viii) as evidence that knowledge is not factive, it could just as well be taken as evidence that 'swonk' is not synonymous with 'know' and so should fall outside of the cosmopolitan investigation. For Hazlett, then, the problem with taking into account cross-linguistic data is that any interesting findings that we take from the data could just as well serve as reason to exclude that data. That much seems correct. What is less clear is why Hazlett thinks this is a fatal problem for the cosmopolitan approach. Instead, Hazlett has just outlined another theoretical possibility that any good theorist must keep in mind. In deciding whether or not 'swonk' should be included as one of the knowledge terms that is the object of study, the theorist must consider how taxonomically disruptive it is to exclude or include that term. If 'swonk' possesses a wide range of properties (semantic or otherwise) that knowledge terms in other languages possess, then it is less taxonomically disruptive to include the term within our class of knowledge terms. But if we find more and more differences between 'swonk' and other knowledge terms, the option of excluding the term becomes more and more plausible. Far from rendering the cosmopolitan project fatally problematic, this is a general situation that a theorist finds themselves in whenever they are employing a taxonomy.¹⁰ Hazlett's objection to the cosmopolitan approach, then, is unconvincing, and so the project of investigating knowledge by employing cross-linguistic data remains a plausible one.

What about the parochialist horn? According to this approach, we should only draw on the intuitions of a single language. Hazlett takes this approach to only be plausible on the assumption that knowledge is a *social kind*, that is its nature and scope supervenes on the social practices within a particular society:

Knowledge, you might argue, is a *social construction*, whose nature and scope supervenes on social practices and institutions [...] Knowledge is more like a status in a rule-governed game, like *being onside* (in association football), than it is a chemical substance like *gold* [...] If knowledge is relative to particular societies, so the argument might continue, it makes sense for us to treat the linguistic practices of English speakers as evidence for and against particular views about the nature and scope of knowledge. For those linguistic practices are a part, perhaps a large or important part, of the social practices and institutions on which knowledge supervenes. (2018, p. 254)

The problem with this approach, according to Hazlett, is that this would not do justice to the idea that philosophy aims to answer questions of a 'broader or a more general nature':

There is a worry here that philosophy cannot be self-consciously parochial, that it must seek to answer questions that are broader and more general. And there is also a worry about method. The empirical and historical methods of sociology, anthropology, and psychology seem more suited to answering questions about the nature and scope of particular social constructions than do the reflective and conceptual methods of philosophy. Of course, philosophy need not be understood in this

¹⁰Thanks to Gabe Dupre for emphasising this point.

way: perhaps philosophy should focus on the particular, in its social and historical context, using suitable empirical methods. But this would raise a pressing question: what does philosophy, so understood, have to *add* to the knowledge generated by sociology, anthropology, and psychology? In any event, my principle worry about the parochialism in the theory of knowledge is that it jettisons that distinctively philosophical interest in things relatively universal: the philosophical question concerns the nature and scope of knowledge, or perhaps human knowledge, but in any event not merely Anglo knowledge. (2018, p. 255)

According to Hazlett, then, the parochialist picture whereby we focus on the linguistic intuitions of a single language, seems to reduce philosophical inquiry to sociological inquiry, in a way that does not square up to the methodology employed by philosophers.

Hazlett means something quite particular by ‘social kind’. He considers a social kind as a kind ‘whose nature and scope supervene on social practices and institutions’ (p. 254). Social kinds may be particular to specific societies, such as (to use Hazlett’s example) the *Burymans’ parade* which takes place in South Queensferry in Scotland. The roles, events and outcomes of this parade are all social kinds; they cannot be understood without appealing to the social practices of South Queensferry. But some social kinds do appear in a number of different societies. Marriage, for example, is a social kind that is in a large number of societies. However, when social kinds do appear in more than one society, one can expect smaller differences regarding the rules, customs and norms surrounding that social kind, as is the case with marriage.

Hazlett is wrong, however, to suggest that the parochial project only makes sense if knowledge is understood as a social kind in his favoured sense. Some kind terms may exist only within certain linguistic communities even when the kind referred to does not supervene on the social practices and institutions of those communities. To illustrate this, consider ‘hygge’. ‘Hygge’ is a Danish term that refers to a particular kind of enjoyable cosiness that can be achieved by sitting by a warm fire, reading a book or having a relaxed lunch with friends. We could imagine hygge as a social kind in Hazlett’s sense, in which case in order to achieve a state of hygge, we might need to be in Denmark and do typically Danish things, perhaps on a particular day of the year and so forth. But this is not how the term is used, and so the term is not a social kind in Hazlett’s sense. Anyone can achieve a state of hygge, whether or not they have the term for it, and whether or not they are in Denmark. Still, the concept is particular to the Danish language, and so to explore the concept further, we really need to speak to Danish speakers to get a better idea of what it involves. In that sense, an analysis of ‘hygge’ would need to be a parochial one.

We can imagine a similar situation with knowledge. To motivate a parochialist analysis of knowledge, we do not need to view knowledge as supervening on social practices or institutions. Nor need it be the case that if we are investigating the English word ‘know’, that we are investigating a kind of state that only occurs in English-speaking countries. For instance, imagine that knowledge as spoken about in English is justified, true belief. Imagine also that no other language has a term that is completely synonymous with the English ‘knows’. Would this mean that justified true beliefs only occur in English-speaking countries? Obviously not. Although other languages would not have a single term that straightforwardly picks out justified true belief in the way that ‘know’ does, speakers of those languages can have beliefs that are justified and true all the same.

Hazlett seems to completely discount this possibility, which is odd as I think that this is a plausible conception of the epistemological project of theorising about knowledge. In this respect, the study of Anglo knowledge can be a study with universal scope insofar as Anglo-knowing is something that one is able to do regardless of the language they speak.

It might be thought at this point that this response to Hazlett’s argument contradicts or is in tension with the response given earlier to Sosa regarding our ability to differentiate between non-linguistic and linguistic affairs. After all, given what I have said about knowledge and

about *hygge*, there is a divide between being in the state of *hygge* and understanding and using ‘*hygge*’—the latter requires some understanding of Danish but the former does not. So does this not thereby show that application of the term is a different matter to whether one can be in the state?¹¹ We need to be careful here to distinguish between a given state being present and the theorist’s project of becoming aware of when that state is present. My point in response to Hazlett concerns the former issue—being in a knowledge state does not require that one possesses the knowledge concept. But being in a knowledge state and being aware that one is in a knowledge state are distinct matters. Our knowledge is not transparent to us, not least because of its factivity, but also possibly due to other externalist conditions on its presence.¹² So when we turn to the theorist’s question of capturing when knowledge is present and when it is absent, in both first-person and third-person cases, what we rely upon is not some guaranteed insight into the full extent of our own knowledge, but a set of intuitions that, as I argued against Sosa earlier, appear to align so closely with our linguistic intuitions so as to be indistinguishable from them.

Both horns of Hazlett’s dilemma have been rejected, and so we no longer have reason to doubt either the parochialist or the cosmopolitan conception of conceptual analysis. But which conception is correct? The answer is: both. There is more than one conceptual analysis project that is of philosophical value. First, there is the parochialist project whereby we can investigate the community-specific classifier adopted by a specific linguistic community. We can investigate knowledge as spoken about in English. In doing so, we are in one sense not investigating an aspect of English society insofar as knowledge will be something possessed across a range of different cultures. But in another sense, we are investigating an aspect of English society insofar as we are investigating a concept that is in circulation in English-speaking communities.

The parochialist project can and should be repeated across a range of different languages. This then opens up the possibility of a cosmopolitan project, whereby we explore the extent to which the concepts employed in various languages match up with one another. They may match up very well, and indeed when questions of translation arise, we often do assume that an expression in one language will be completely synonymous and co-extensive with an expression in another. It will not be surprising, however, if closer inspection reveals that this assumption is often false. In the case of ‘know’, it is tempting to think that expressions in other languages (e.g., ‘savoir’, ‘kennt’, ‘conoscere’) are completely synonymous with ‘know’. Perhaps, if Wierzbicka (2018) and others within the *Natural Semantic Metalanguage* approach are correct, then KNOW is a universal concept that is present in *every* language. But this seems like a substantive empirical claim and we should be prepared for the possibility that there is no straightforward relationship of synonymy across the knowledge terms of various languages. There may even be subtle differences that are hard to detect—it may be that the knowledge term of one language may require a slightly higher standard of justification than other languages. If that were so, it would be a difficult, but not impossible, difference to tease out via investigation.

The cosmopolitan project clearly depends on the parochial project insofar as cross-linguistic analysis will depend on an analysis in each of the languages appealed to. But it may be that the parochial project will also be able to draw upon the cosmopolitan. In particular, if a set of related terms (e.g., knowledge terms) across languages are identified, and stable patterns are identified across those languages (e.g., whenever a knowledge term in a language has syntactic property α , it tends to have semantic property β), then this can help inform the parochial project for new languages by giving us patterns that we can expect to be repeated. Or to use a more familiar example, if English were the only language that has a knowledge term that allows for Gettier cases, then this would provide defeasible evidence to look again at the initial Gettier

¹¹Thanks to an anonymous referee for pressing this point.

¹²Of course, if Williamson (2002) is right, no non-trivial mental state is transparent.

cases and consider whether the intuitive evidence is as we thought it was, or whether some alternative explanation is available that would then allow English to fall in line with other languages with regard to its knowledge terms.

Much of the previous epistemological work, I contend, can be viewed through this parochialist lens. But this leads to a natural worry. We may start an epistemology course by asking students to read Plato's *Meno* or Descartes' *Meditations*. We freely draw upon texts originally written in different languages all the time in philosophy. If conceptual analysis is to be understood through a parochialist lens, are we in error when we draw upon such texts? Are we confusing one parochialist project for another? This, however, is too quick. The value of such texts for us is that when translated into *our* language, the arguments and considerations contained in the text reveal truths to us about the way that we use our knowledge terms. When we read that Socrates is concerned to explain the fact that knowledge is more valuable than mere true belief, we see that this is something that requires explanation according to the terms in our own language. Whatever relation it is that we refer to by the term 'know', it seems to be more valuable than the relation we refer to by the term 'truly believe'. Our agreement with a translated Socratic text helps guide our parochialist project. Alternatively, the value in reading such texts can lie precisely in the fact that concepts not lexicalised in our language are introduced to us.¹³

Still, this may lead to a yet further worry. Once we acknowledge that the conceptual analysis project is (in the first instance) *parochialist* and that the object of our study is, for example, knowledge as spoken of by English speakers, it can be asked: *who cares about that?* Even if we note, as we did earlier, the fact that 'know' as used in English applies to individuals inside and outside of English-speaking communities, once we are faced with a potentially wide array of knowledge terms across a number of languages and with varying meanings, then the particular state that is picked out by the English term, differing as it may do with the states picked out by the corresponding terms in other languages, is just one epistemic state among many. Does this not minimise the importance of the philosophical project?

One should generally be wary of objections based on whether a given project is interesting or not. That said, defending the significance of the parochialist project should not be difficult, provided that the epistemic terms in question do play a significant role within that community. And I take it as fairly uncontroversial that, at least in the case, of 'know', there is some significant role played. This is reflected in the philosophical literature, where it has variously been suggested that knowledge is a source of praise and criticism, that it serves to affect both practical action and inquiry, and that it serves a social function insofar as it helps us identify others as reliable sources of information. Whether or not the particular theories associated with these more general claims are correct, I take it that they are all suggestive of a less controversial claim: that 'know' plays a role of significance within our epistemic lives. The importance of the knowledge concept is also reflected in its sheer prevalence. As Hansen et al. (2021, p. 247) note, 'know' is one of the most frequent verbs in English even once discourse marking uses are discounted. This should be sufficient to warrant a philosophical investigation into its nature.

Still, even if it is conceded that the view of conceptual analysis defended here would indeed be of philosophical significance, it could nevertheless be argued that it will fail to capture the *evaluative* and *normative* side of epistemology. Stephen Stich has done more than anyone to outline an objection of this kind, and so it is to his work that I now turn.

¹³For a clear example, see Pasnau (2013) for an account of Aristotle's conception of *epistēmē* and Descartes' notion of *scientia* as ideal epistemic notions with no straightforward translation in English.

5 | STICH ON EPISTEMIC NORMATIVITY AND COGNITIVE DIVERSITY

Epistemology is a normative discipline in that one epistemological goal is to capture not only the belief-forming practices that we in fact use, but in some sense the belief-forming practices that we ought to use. Stich has argued on a number of occasions (Stich, 1988, 2009; Weinberg et al., 2001) that an approach to epistemology centred on conceptual analysis cannot achieve that goal. Stich's arguments have been responded to by Sosa (2009) and I take those responses to be important. However, I will argue that the debate between Sosa and Stich ignores the theoretical possibility that an account of our epistemic concepts could come apart from an account of the epistemic normative facts, even while acknowledging important links between the two. This possibility must be ruled out before Stich's objection can have force. So in this section, I will first survey the debate between Stich and Sosa, before outlining this ignored possibility.

Stich (1988) initially raises this objection in response to Goldman's (1986) conceptual analysis approach to justification.¹⁴ Following Goldman, he notes that one of the central goals of a theory of justification is to produce a set of 'J-rules' that specify epistemically permissible ways to form a belief. In producing a theory, we need a criterion of rightness that rule certain J-rule sets out and rule certain J-rule sets in. For the conceptual analysis approach Stich considers, the criterion of rightness, will be that the J-rule set aligns with the ordinary conception of justification as reflected in ordinary language and ordinary thought. But he argues that this approach cannot capture the normative element of epistemology, given the possibility that the practices of epistemic evaluation reflected in ordinary thought and ordinary language may vary across different cultures:

[I]magine that we have located some exotic culture that does in fact exploit cognitive processes very different from our own and that the notions of epistemic evaluation embedded in their language also differ from ours. Suppose further that the cognitive processes prevailing in that culture accord quite well with *their* evaluation notions, while the cognitive processes prevailing in our culture accord quite well with *ours*. Would any of this be of any help at all in deciding which cognitive processes we should use? Without some reason to think that one set of evaluative notions was preferable to the other, it seems clear that it would be of no help at all. (Stich, 1988, p. 407)

The very possibility of such cultural variation reveals that conceptual analysis understood as capturing the community-specific classifiers circulating within a linguistic community must seemingly remain silent on whether the classifiers we are using are the classifiers we ought to be using.

As Sosa (2009) has noted in response, there is no principled barrier here to a kind of pluralism regarding epistemic value. If an analysis of one community's epistemic practices reveals that they are valuing one kind of epistemic commodity (to use Sosa's phrase), while the epistemic practices of another community reveal that they are valuing another kind of epistemic commodity, then this just reveals that there is more than one epistemic commodity that can be valued:

The fact that we value one commodity, called 'knowledge' or 'justification' among us, is no obstacle to our also valuing a different commodity, valued by some other community under that same label. And it is also compatible with our learning to

¹⁴Elsewhere, I have criticised the idea that epistemic justification is a suitable candidate for conceptual analysis. If there is some common notion of epistemic justification, it is not plausibly investigated by considering the ordinary usage of 'justify'. See Grindrod (2022).

value that second commodity once we are brought to understand it, even if we previously had no opinion on the matter. (Sosa, 2009, p. 109).

In a further response to Sosa, Stich (2009, p. 235) has argued that the kind of pluralism advocated by Sosa here is problematic given that epistemology understood as a normative discipline should also provide *rules of permissibility* so that it can reflect the fact that there are right and wrong ways to behave as an epistemic agent. In the case of belief-formation, it should provide rules stating when it would and would not be permissible to form a belief. Two cultures with distinct epistemic practices would (or at least could) be following distinct permissibility rules that would conflict with one another. For instance, we can imagine two different cultures that differ according to their respective belief-norms, such that one follows (ix) and the other follows (x):

- ix. One ought to believe that p iff doing so will result in anti-Gettiered JTB that p.
- x. One ought to believe that p iff doing so will result in JTB that p.

These two rules provide conflicting verdicts in Gettier cases, and so Stich's point is that it is incumbent upon an epistemology to tell us which norm (if either) we ought to adopt.

In defence of Sosa's position, Ichikawa (2014, pp. 247–248) responds by emphasising that norms regarding when we ought to believe are only plausible if they have *ceteris paribus* clauses, in pain of being too demanding. As Ichikawa puts it, without these clauses they would 'prohibit building houses in favour of counting bricks' (Ichikawa, 2014, p. 248). And once such principles are qualified in this way, the conflict between, for example, (ix) and (x) is far less clear. It may be instead that cases such as Gettier cases are the kinds of difficult cases that the *ceteris paribus* clause allows the norm in question to be silent on.

An alternative response is suggested in Sosa's (2009, p. 110) claim that epistemic normativity is restricted to its domain, such that epistemic norms are not to be interpreted as placing restrictions on our behaviour all things considered, but only within the epistemic realm. Sosa uses the analogy of the norms and evaluations around shooting, but we can also look to games. We can, for instance, imagine a game according to which the rules of play regarding some particular moment in the game is best described via a biconditional. Still, the path from this biconditional within the game's domain to what we ought to do, all things considered, is complicated by weighing up a myriad of other factors. For instance, whether I should make a move in the game at all is affected by whether there is an uncontrollable fire in the room I am in. Still, this does not affect the truth of the biconditional in question.

How does this affect Stich's original concern about cognitive diversity? Recall that Stich has in mind two communities with distinct epistemic practices, such that 'know' has different extensions across the two communities, the two communities value different epistemic commodities, and there are different epistemic norms of permissibility associated with each community. Sosa's suggestion is that we can treat the two distinct communities as two distinct epistemic domains, so that no conflict between the permissibility rules arise:

[T]here seems no more reason to postulate such conflict than there would be when we compare someone who rates cars in respect of how economical they are with someone who rates them in respect of how fast they go. (Sosa, 2009, p. 110)

On this picture then, different communities may have different epistemic domains, complete with their own internal norms of permissibility. Such norms do not conflict with one another, and I can choose to engage with more than one domain provided that I value the commodities of each, in which case my epistemic behaviour could be evaluated from the standpoint of either. But while we have carefully avoided the potential conflict that Stich raises, something like his

original concern still remains. That is, if we expected our final epistemology to tell us which epistemic practices we ought to be engaging in, including which epistemic norms we ought to be following, then Sosa's pluralism does seem to remain silent in this regard. This dissatisfaction may be more acute when we consider the possibility of epistemic communities with dramatically different epistemic practices to our own. Communities that differ over Gettier cases is one thing, but we can at least imagine communities that value refraining from belief or that value holding false beliefs, or beliefs that can be stated quickly and so forth. Should our epistemology not rule these communities as having gone wrong somewhere?

Sosa does, however, suggest an answer to this question in seeking to account for what he imagines as a disagreement in epistemic practices between a community that employs seemingly reasonable epistemic practices (what he calls the *enlightened community*) and a community that employs seemingly questionable epistemic practices (what he calls the *superstitious community*):

What the enlightened object to is the notion that the sort of status elevated by the superstitious constitutes *epistemic* value in the actual world. And this is presumably because they see superstitious status as insufficiently connected with truth.

[...]

Epistemic justification concerns specifically *epistemic* values, such as truth, surely, and perhaps others not entirely reducible to truth, such as understanding (Sosa, 2009, pp. 110–111)

Here Sosa shows how differing verdicts between the enlightened community and the superstitious community can be drawn. The practices of the superstitious, while they may prove valuable within some other domain, could not be deemed epistemically valuable because they are 'insufficiently connected with truth'. Here we find that truth plays a fundamental role in the epistemic realm insofar as whatever is epistemically valuable must be understandable in terms of *some* connection to truth. It may be because such commodities are instrumental to or reducible to true belief, or it may be that there is some more indirect connection with truth. While the kinds of possible connections a given commodity might have with truth in order to count as an epistemic commodity requires further development, we at least have a starting point with which to rule out certain practices as epistemically problematic, such that we do have a way of adjudicating between on the one hand, our own epistemic practices, and on the other, that of the superstitious and other communities with exotic belief-forming practices.

In this respect, then, there is an important restriction on Sosa's pluralism. While there may be numerous epistemic commodities and numerous epistemic norms that arise from them, what can count as an epistemic commodity is limited. For the purposes of this paper, the important point is that Sosa's pluralism is one way in which the view of conceptual analysis that has been defended in this paper can respond to Stich's worry regarding cognitive diversity.

It is also notable, however, that some will find this kind of pluralism unpalatable insofar as it is not clear whether it is consistent with a value monism and in particular with *veritism*: the view that (all and) only true beliefs are epistemically valuable and (all and) only false beliefs are epistemically disvaluable. Veritists argue that truth is the only source of epistemic value and that any other epistemically valuable entities are only instrumentally valuable insofar as they are a means to truth. I want to finish this paper by considering a way of responding to Stich's concern while remaining neutral on the truth of veritism or other monist views of epistemic value. In doing so, I intend to better spell out how the normative strand of epistemology should be understood given the view of conceptual analysis defended here.

To understand this alternative approach, it is useful to distinguish between the project of providing a conceptual analysis of the epistemic notions within a community and the project of

understanding the evaluative and normative dimensions within the epistemic realm. Stich's objection arises because it is assumed that if conceptual analysis is to be feasible then it must provide us with a means to the latter as well as the former. In considering two different communities that may differ in what they mean by 'know', for instance, the thought was that this may reflect that fact that the two communities differ in what is epistemically valuable. But if we take the view of conceptual analysis that I have defended as our starting point, then whether evaluative notions such as 'know' do serve to perfectly track what is epistemically valuable in this manner is a substantive question. It may be that if two societies differ in what they are willing to apply their closest translation of 'know' to, then this reflects that they value two distinct epistemic commodities (as Sosa suggests). Alternatively, it may be that the two societies are alike in what they value epistemically, but they differ in the epistemic tools they use.

Why would two communities differ in their epistemic tools given that they value the same commodities? To see why, let us consider a somewhat artificial comparison of two imaginary communities. Suppose that these two communities have very similar knowledge concepts: society A requires that the subjective probability of a belief exceeds 0.7 while society B requires that the subjective probability of a belief exceeds 0.9. Assume for the sake of argument that a form of veritism is true across all communities—so only true beliefs are valuable and only false beliefs are disvaluable. We may ask of these societies, which is better set to attain as many true beliefs as possible while also avoiding false beliefs, as according to veritism, it will be this society that is employing the superior notion. But the answer to this question will depend partly on the way the epistemic environment is, in particular on: (a) how plentiful evidence is in the world; and (b) how surprising or deceptive the world is, that is how misleading the available evidence is. To take one limiting case where the world is in fact an evil demon world, neither society is particularly successful at attaining true beliefs, and in fact their knowledge-concepts would best be scrapped in favour of a maximum threshold concept so that they at least form no false beliefs. In less surprising worlds, where the evidence is far less misleading, and so evidence of 0.7 probability will suffice more often than not, the less conservative society A will have the superior epistemic practices and form more true beliefs. But in more surprising worlds, society B will have the superior epistemic practices.

All this is to say that even once we hold fixed a theory of epistemic value, what the optimal form of epistemic practice is going to be will depend in part on the way the world is. This is sufficient to show that two communities may admit of differing analyses of their epistemic classifiers even while a single theory of epistemic value remains true for both. This also reveals that this is not a fixed issue—that our epistemic classifiers that form part of our epistemic practices may evolve over time in response to a changing environment, because the *optimal* set of epistemic practices may change over time.

Furthermore, it would not be surprising if the actual formation of our epistemic classifiers, as well as the manner in which they evolve over time, are themselves affected by *non-epistemic* factors. For instance, in a world where false positives (i.e., the formation of false beliefs) are much more costly than false negatives (failure to form true beliefs) it may well be that our epistemic notions naturally evolve to track that, by becoming more conservative.¹⁵ Again, this is reason to think that two societies could differ in their epistemic classifiers even while a single, monist theory of epistemic value remains true for both.

Finally, even if the epistemic environment remains the same across two cultures, it would not be particularly surprising, I suggest, if the epistemic classifiers that enjoy common currency only *imperfectly serve* as tools with which to maximise epistemic value. A classifier can develop over time, not only because it is reacting to the changing epistemic landscape, but because it is

¹⁵Note that this would not constitute a form of pragmatic encroachment. Pragmatic encroachment states that our epistemic classifier 'know' is such that whether it is applicable to a given subject regarding a given proposition is dependent upon certain practical factors. I am not making that claim here. I am claiming that which epistemic classifier ends up being circulated within a given community could depend on practical factors.

subject to particular pressures in virtue of being in common currency. Simplicity is one such pressure, as in order for a classifier to enjoy wider circulation, it must be that it is relatively simple to employ.

The above considerations are all reasons to think that the relationship between a conceptual analysis of epistemic classifiers and a theory of epistemic normativity is not as straightforward as is seemingly assumed in the debate between Sosa and Stich. For this reason, it is important to be clear that to the extent that epistemology is a normative discipline, it is not plausible that it can start and end with conceptual analysis. Instead, a theory of epistemic normativity should at first be developed independently, so that we have some account of what is epistemically valuable and what the correct and incorrect ways to behave epistemically are.¹⁶ This then provides us with a method of evaluation regarding the epistemic classifiers we in fact use. If all is well, our epistemic classifiers are the optimal tools required in order to behave in the best epistemic way that we can, but we can also be open to the alternative possibility that alternative classifiers would better fulfil the role. It is at this point that a case for conceptual engineering can be given, so as to improve upon the classifiers that circulate within a community.

To summarise this section, Stich has argued that conceptual analysis within epistemology will prove unable to capture the normative and evaluative elements of epistemology, particularly with regard to evaluating whether the epistemic practices of a given community are good practices, or even the right practices. Sosa has provided an important response to this worry, by defending a form of pluralism such that we could view each epistemic community as a distinct epistemic domain, provided that the commodities valued by each community are sufficiently connected with truth. I have also given a number of reasons to think that epistemic classifiers may only imperfectly reflect what is valued within an epistemic community, so that it is best to treat normative epistemology, initially at least, as a distinct project from the conceptual analysis of our epistemic classifiers.

6 | CONCLUSION

I have developed and defended a view of conceptual analysis within epistemology as a linguistic endeavour. According to that view, we must be explicit that in providing a conceptual analysis of, for example, knowledge, we are targeting the concept associated with the English term 'know' and its cognates, and so our conclusions are in the first instance parochialist. Contrary to Kukla, there is not sufficient reason to think that the project is endangered by the supposed messiness of epistemic discourse, nor of the use of knowledge terms in academic disciplines. Contrary to Hazlett, the project *can* be thought of as capturing the universal nature of the epistemology (and philosophical inquiry more generally) insofar as knowing (as spoken of in English) is something that anyone across any community can achieve. And contrary to Stich, the project is entirely consistent with a full theory of epistemic normativity, even if conceptual analysis will not be the route by which we obtain that theory.

ACKNOWLEDGEMENTS

This work was previously presented to the Reading philosophy graduate class. Thanks to Nat Hansen for comments on a much earlier version of this paper.

¹⁶I will not develop a methodology for normative epistemology here. But I will note that while it does require something other than a conceptual analysis of our epistemic classifiers, it may still be intuition-based. After all, a range of intuitions regarding epistemic value and normativity are already discussed in the literature (e.g. regarding epistemic conduct, regarding whether it is epistemically valuable to form useless beliefs, regarding credibility judgements, etc.), and these are often taken as the basis with which to establish further normative claims.

ORCID

Jumbly Grindrod  <https://orcid.org/0000-0001-8684-974X>

REFERENCES

- Basu, R. (2019) Radical moral encroachment: the moral stakes of racist beliefs. *Philosophical Issues*, 29(1), 9–23. Available from: <https://doi.org/10.1111/phils.12137>
- Beebe, J.R. & Buckwalter, W. (2010) The epistemic side-effect effect. *Mind & Language*, 25(4), 474–498. Available from: <https://doi.org/10.1111/j.1468-0017.2010.01398.x>
- Brown, J. (2011) Thought experiments, intuitions and philosophical evidence. *Dialectica*, 65(4), 493–516.
- Cappelen, H. (2012) *Philosophy without intuitions*. Oxford: Oxford University Press.
- Chudnoff, E. (2017) The reality of the intuitive. *Inquiry: An Interdisciplinary Journal of Philosophy*, 60(4), 371–385.
- Cummins, R.C. (1998) Reflections on reflective equilibrium. In: DePaul, M. & Ramsey, W. (Eds.) *Rethinking intuition*. Lanham, MD: Rowman & Littlefield.
- Deusch, M. (2010) Intuitions, counter-examples, and experimental philosophy. *Review of Philosophy and Psychology*, 1(3), 447–460. Available from: <https://doi.org/10.1007/s13164-010-0033-0>
- Fritz, J. (2017) Pragmatic encroachment and moral encroachment. *Pacific Philosophical Quarterly*, 98(S1), 643–661.
- Goldman, A.I. (1986) *Epistemology and cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A.I. (2010) Philosophical naturalism and intuitional methodology. *Proceedings of the American Philosophical Association*, 84(2), 115–150.
- Goldman, A.I. & Pust, J. (1998) Philosophical theory and intuitional evidence. In: *Rethinking intuition*. Lanham, MD: Rowman & Littlefield.
- Grice, H.P. (1989) Postwar Oxford philosophy. In: *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Grindrod, J. (2022) Justification: insights from corpora. *Episteme*, 1–25. Available from: <https://doi.org/10.1017/epi.2022.44>
- Hansen, N., Porter, J.D. & Francis, K. (2021) A corpus study of ‘know’: on the verification of philosophers’ frequency claims about language’. *Episteme*, 18(2), 242–268. Available from: <https://doi.org/10.1017/epi.2019.15>
- Hazlett, A. (2010) The myth of factive verbs. *Philosophy and Phenomenological Research*, 80(3), 497–522.
- Hazlett, A. (2018) Theory of knowledge without (comparative) linguistics. In: Stich, S., Mizumoto, M. & McCreedy, E. (Eds.) *Epistemology for the rest of the world*. Oxford: Oxford University Press.
- Horne, Z. & Cimpian, A. (2018) Subtle syntactic cues affect intuitions about knowledge: methodological and theoretical implications for epistemology. In: Lombrozo, T., Knobe, J. & Nichols, S. (Eds.) *Oxford studies in experimental philosophy*, Vol. 2. Oxford: Oxford University Press.
- Ichikawa, J.J. (2014) Who needs intuitions? Two experimentalist critiques. In: Booth, A.R. & Rowbottom, D.P. (Eds.) *Intuitions*. Oxford: Oxford University Press.
- Kim, M. & Yuan, Y. (2015) No cross-cultural differences in the Gettier car case intuition: a replication study of Weinberg et al. 2001. *Episteme*, 12(3), 355–361. Available from: <https://doi.org/10.1017/epi.2015.17>
- Kukla, Q. (2015) Delimiting the proper scope of epistemology. *Philosophical Perspectives*, 29(1), 202–216. Available from: <https://doi.org/10.1111/phpe.12062>
- Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N. et al. (2015) Gettier across cultures. *Noûs*, 51(3), 645–664.
- Nado, J. (2016) The intuition deniers. *Philosophical Studies*, 173(3), 781–800.
- Pasnau, R. (2013) Epistemology idealized. *Mind*, 122(488), 987–1021. Available from: <https://doi.org/10.1093/mind/fzt093>
- Schaffer, J. (2007) Knowing the answer. *Philosophy and Phenomenological Research*, 75(2), 383–403. Available from: <https://doi.org/10.1111/j.1933-1592.2007.00081.x>
- Seyedsayamdost, H. (2015) On normativity and epistemic intuitions: failure of replication. *Episteme*, 12(1), 95–116. Available from: <https://doi.org/10.1017/epi.2014.27>
- Sosa, E. (2009) A defense of the use of intuitions in philosophy. In: Murphy, D. & Bishop, M. (Eds.) *Stich and his critics*. Oxford: Blackwell, pp. 101–112.
- Stich, S. (1988) Reflective equilibrium, analytic epistemology and the problem of cognitive diversity. *Synthese*, 74(3), 391–413.
- Stich, S. (2009) Replies. In: Murphy, D. & Bishop, M. (Eds.) *Stich: and his critics*. Oxford: Blackwell, pp. 190–252.
- Swain, S., Alexander, J. & Weinberg, J.M. (2008) The instability of philosophical intuitions: running hot and cold on Truetemp. *Philosophy and Phenomenological Research*, 76(1), 138–155.
- Weinberg, J.M. (2017) What is negative experimental philosophy good for? In: D’Oro, G. & Overgaard, S.E. (Eds.) *The Cambridge companion to philosophical methodology*. Cambridge: Cambridge University Press, pp. 161–184.
- Weinberg, J.M., Nichols, S. & Stich, S. (2001) Normativity and epistemic intuitions. *Philosophical Topics*, 29(1 & 2), 429–460.

- Wierzbicka, A. (2018) I know: a human universal. In: Stich, S., Mizumoto, M. & McCready, E. (Eds.) *Epistemology for the rest of the world*. Oxford: Oxford University Press.
- Williamson, T. (2007) *The philosophy of philosophy*. Malden, MA: Blackwell Publishing.
- Williamson, T. (2002) *Knowledge and its limits*. Oxford: Oxford University Press.

How to cite this article: Grindrod, J. (2023) A defence of conceptual analysis as a linguistic endeavour. *Theoria*, 1–19. Available from: <https://doi.org/10.1111/theo.12477>