



Inquiry

An Interdisciplinary Journal of Philosophy

ISSN: (Print) (Online) Journal homepage: <https://www.tandfonline.com/loi/sinq20>

Re-engineering knowledge: a case study in pluralist conceptual engineering

Jennifer Nado

To cite this article: Jennifer Nado (2021): Re-engineering knowledge: a case study in pluralist conceptual engineering, *Inquiry*, DOI: [10.1080/0020174X.2021.1903987](https://doi.org/10.1080/0020174X.2021.1903987)

To link to this article: <https://doi.org/10.1080/0020174X.2021.1903987>



Published online: 30 Mar 2021.



Submit your article to this journal [↗](#)



Article views: 63




View related articles [↗](#)



View Crossmark data [↗](#)



Re-engineering knowledge: a case study in pluralist conceptual engineering

Jennifer Nado 

Department of Philosophy, University of Hong Kong, Pokfulam, Hong Kong

ABSTRACT

When fans of conceptual engineering discuss examples of their craft, they frequently focus on cases of ‘one-to-one’ conceptual engineering. That is to say, they focus on cases where a target pre-engineering concept is revised into, or replaced by, a single successor. The possibility that we might instead replace a suboptimal concept with *multiple* successors is, by contrast, comparatively underexplored. The goal of this paper is to defend this type of pluralist conceptual engineering as legitimate, and as a promising approach to unravelling certain traditional philosophical puzzles. I’ll do so largely by way of example, by exploring a preliminary application of the pluralist approach to one of the most notoriously recalcitrant targets of conceptual analysis: knowledge.

ARTICLE HISTORY Received 1 December 2020; Accepted 31 December 2020

KEYWORDS Conceptual engineering; metaphilosophy; pluralism; contextualism; pragmatic encroachment; knowledge norms

When fans of conceptual engineering discuss examples of their craft, they frequently focus on cases of ‘one-to-one’ conceptual engineering. That is to say, they focus on cases where a target pre-engineering concept is revised into, or replaced by, a single successor. The possibility that we might instead replace a suboptimal concept with *multiple* successors is, by contrast, comparatively underexplored.¹ Yet ‘one-to-many’ conceptual engineering seems like a natural strategy in certain cases. Most obviously, such a strategy might be employed when the pre-engineering concept conflates categories that we hold to be genuinely distinct, as in for instance the move to recognize the difference between biological sex and gender.

If one is drawn to a functional or practical approach to conceptual engineering, then one-to-many conceptual engineering becomes even

CONTACT Jennifer Nado  jennifernado@gmail.com

¹A notable exception here is Kevin Scharp’s excellent work on replacing the pre-theoretic concept of truth (see Scharp (2013)).

more attractive. On such views, the goal of conceptual engineering is to improve the target concept's ability to serve its function, fulfill its purpose, play its role, or something of the like. But plausibly, many of our ordinary pre-engineering concepts are 'multitaskers', playing multiple roles or serving multiple functions. And a concept that is a 'jack-of-all-trades' is, arguably, likely to be a master of none. If this is correct, then we may find our purposes better served by a plurality of concepts, each custom-designed to best fill one of the original concept's roles.

The goal of this paper is to defend this type of engineered 'division of labour' as legitimate, and as a promising approach to unravelling certain traditional philosophical puzzles. I'll do so largely by way of example, by exploring a preliminary application of the pluralist approach to one of the most notoriously recalcitrant targets of conceptual analysis: knowledge. In line with a functional/practical approach to engineering, a central focus of the paper will be an examination of the various purposes which epistemologists have suggested our current knowledge-concept may serve. The very diversity of such purposes already suggests that our current concept may be a bit of a 'kludge', fulfilling multiple purposes in a passable but non-optimal fashion – and that, at least for philosophical purposes, a set of specialized, technical successor concepts might better serve. Exploring features that would improve a concept's suitability to play each of these roles, I'll argue, raises the serious possibility that no *single* concept can offer ideal performance under such a workload.

1. Methodological background

Readers of this special issue are likely already possessed of at least a passing acquaintance with the nature and argued merits of conceptual engineering as a general method. Rather than re-tread that terrain, I suggest we jump right into particulars. My primary aim in this paper is to illustrate the virtues of a pluralist approach by *engaging* in conceptual engineering (albeit of a *very* preliminary and speculative sort); but before we start to get our hands dirty, we will need some background on the particular flavor of conceptual engineering that will underwrite my defence of subdividing our knowledge-concept. I'll also delve a bit into the methodological implications of said flavor, and how they vindicate pluralist proposals as legitimate instances of conceptual engineering.

I've used the words 'functional' and 'practical' as labels for the approach to conceptual engineering that I endorse – and since 'practical' is perhaps the more theoretically neutral of the two terms, I'll stick with

that characterization from here on out. A ‘practical’ approach to conceptual engineering is a general methodological orientation that naturally contrasts with what we might call the ‘semantic’ approach. A semantic approach to conceptual engineering typically focuses on what might be seen as ‘semantic defects’ – is the current concept vague, or internally inconsistent, or empty, or in some way incoherent? Can these semantic defects be remedied? A practical approach, by contrast, focuses on the functions or purposes that the target concept serves, and asks whether the concept can be altered to fulfill those purposes more effectively.²

For a practical conceptual engineer, success in engineering will be primarily measured by how effectively the proposed concept fulfills the purposes it is intended to serve (though of course there may be other desiderata, such as simplicity). As a quick example, suppose (as many of us do) that one purpose of the concept of free will is to help identify potential bearers of moral responsibility. If a practically-minded engineer were to aim to improve this concept, she would then ask: are there any changes to our current free will concept that would enable it to fulfill this purpose more effectively? She might argue, for instance, that incompatibilist characterizations of free will ought to be rejected – not on grounds of falsity, but on the grounds that the actual world is such that incompatibilist notions of free will imply universal lack of moral responsibility, and the ability to attribute moral responsibility is a necessity for a well-run society. Of course, I don’t mean to endorse this particular revisionary argument; but it should give a sense of the sort of strategy a fan of the practical approach to conceptual engineering would apply.

There are a number of questions one might ask about the details of the practical approach. Answering said questions in any detail is beyond the scope of this paper – we’ve got engineering to do, after all. But I’ll give quick indications of my own preferred stances on two questions of particular importance.

Question: What exactly is the function or purpose of a concept?

Many proponents of practical approaches to conceptual engineering frame their views in terms of functions, which immediately raises all sorts of questions regarding the nature of functions, how concepts could have such things, how the function of a concept might be identified, and so forth. My preference is to shy away from the notion of ‘function’ in favor of appeal to the *purposes* for which a concept is used. The

²The practical approach has enjoyed substantial popularity in the recent conceptual engineering literature; see for instance Haslanger (2000, 2020), Prinzing (2018), Simion and Kelp (2019), and Thomasson (2020).

concept of marriage, for instance, is used for many purposes: to determine who is eligible to inherit, who has custody over children, who may or may not occupy certain religious roles, and so forth. None of these is plausibly ‘the’ function of the marriage concept, but all of these are relevant to evaluating the success of a conceptual revision. For instance, one might argue in favor of including same-sex partnerships under ‘marriage’ by appeal to the fact that it is morally preferable to allow same-sex partners to possess the above-mentioned rights. ‘Purpose’ in the sense I intend to use it is much broader than ‘function’, even if one is comfortable with the possibility of concepts having multiple functions: consider, e.g. the fact that books are often used for the purpose of stabilizing wobbly table legs, though that is certainly not plausibly one of their *functions*.

In fact, I would be inclined to take an even broader approach, recognizing the relevance of practical factors that aren’t even classifiable as ‘purposes’, such as ease of learnability. On my view, just about any answer to the question ‘why ought we to categorize things *this way*, rather than *that*?’ will be a legitimate argumentative resource which the practical conceptual engineer may use to promote her revisions. This leaves in a lot; answers to the aforementioned question might include ‘because *that* way of sorting has negative social or political consequences’, or ‘because *that* way of sorting is needlessly complex and ad hoc’, or even simply ‘because *this way* better reflects the joints of nature’.

This broad approach to the relevant practical features of a concept is the primary reason why I hold that pluralist engineering proposals are a natural and salutary form of conceptual engineering. Perhaps a ‘pure’ functionalist might aim to identify a single, core, proper function to each of our concepts, and argue that any legitimate successor must preserve this function. On such a view, there’s little reason to retain more than one concept at the end of our tinkering – the one that best serves said function. But on the broader perspective that I advocate, in nearly all cases there will be multiple purposes that a concept is used for, and multiple roles it plays in our cognitive lives. Thus, an obvious strategy is to aim to redistribute the workload, by devising one or more successor concepts which focus only on some subset of the original’s ‘jobs’. Of course, a novel concept which serves only some of the roles of the pre-engineering target will typically not be suited to wholly displace its predecessor. But on the practical approach I advocate, this isn’t problematic. Nothing prevents us from devising additional successors to handle the remaining purposes, or even retaining the original concept for certain uses which it seems to handle suitably well.

As an illustration of the latter scenario, consider an example adapted from Carnap: the explication of the ordinary concept ‘warmth’ into the quantitative concept ‘temperature’. For most scientific purposes, use of ‘temperature’ is superior to use of ‘warmth’ due to the former’s objectivity and precision. But these very features make it less well-suited to everyday, casual reports on the weather. Due to factors like wind chill and humidity, subjective feelings of warmth often come apart from temperature, and it is warmth that most people are interested in before they step out the door. There are good reasons, then, for retaining both forms of measurement. This example illustrates a case of ‘partial’ replacement – in some contexts in which we previously used ‘warm’ we now use a precise temperature measurement, and in others the original concept is retained. But partial replacements are just as pluralist as replacements with multiple bespoke successors; after all, at the end of the process, we now use multiple concepts where previously there was but one.

Question: How does the practical approach deal with ‘Strawson’s Challenge’?

As most readers will likely know, Strawson famously objected to Carnap that explication threatens to ‘change the subject’ (Strawson 1963). The concern underlying Strawson’s worry is that changing the subject will result in important philosophical questions going unanswered. Philosophers want to know about the nature of *knowledge*; if our re-engineering practices are too liberal, however, then at the end of the conceptual engineering process we will no longer be talking about knowledge. We’ll be talking about knowledge*, and no amount of insight into knowledge* will answer our original philosophical questions about knowledge.

Prima facie, this poses a serious challenge to pluralist approaches – if replacing knowledge with knowledge* inappropriately changes the subject, then replacing knowledge with knowledge₁, knowledge₂, knowledge₃ and so forth looks to toss the subject right out the window. But I’d argue that this ‘worry’ is largely illusory. Abandoning our original questions is only problematic, after all, if those questions were worth our time in the first place. So we must ask ourselves: why is *knowledge* held to be so much more philosophically weighty than any of the *other* possible ways of categorizing epistemic states? Why not instead spill our philosophical ink over some more tractable concept, like good old JTB?

If the concept of knowledge is genuinely more worthy of philosophical scrutiny than JTB, I hold that it must be because – unlike JTB – it plays a

certain valuable role or roles in our cognitive lives. It serves a certain set of purposes. It is, in other words, a *useful* way of categorizing. But if a pluralist conceptual engineer does her job well, then her proposed set of successor concepts knowledge₁ through knowledge_n will serve all the purposes the concept knowledge formerly served – and better, to boot. I'd suggest that, as a result, the reasons for philosophizing about knowledge₁ through knowledge_n will wholly displace any reasons we might have had for studying knowledge. The more conservative among us may balk at all this, and protest that knowledge is simply *intrinsically* of philosophical interest. Well, if you simply can't rid yourself of the desire to know the nature of *knowledge*, there's nothing in my approach to conceptual engineering that forbids you from keeping the old concept around in addition to the new. Though I'd suggest some reflection on why that *particular* carving of logical space strikes you as so especially distinguished, if the roles it plays can be better filled by well-fashioned replacements!

2. The many roles of knowledge

That's enough by way of preparatory groundwork; let's get to the main event. On the practical approach to conceptual engineering, the place to begin a revisionary project is to seek at least a preliminary account of the purposes for which the target concept – in our case, knowledge – is used. Fortunately, we don't need to start from scratch here. The epistemological literature is already rife with discussions of various roles that knowledge allegedly plays in our cognitive and social lives. In an excellent recent paper, Fassio and McKenna (2015) have surveyed the most prominent of these in service of their own revisionary account of knowledge. Before reviewing the roles they identify, however, I'd like to make a brief detour to discuss their methodology and how it differs from my own.

Fassio and McKenna make their case for revising the knowledge concept by looking at two separate features of that concept – its plausible purposes, which we will survey shortly, and a set of *principles* which they take to govern its use. The latter includes, for example, the 'Parity of Evidence Principle' (roughly: if two subjects have the same evidence regarding some proposition *p*, then either they are both in a position to know *p* or neither is), and the 'Social Principle' (which asserts that the conditions for correctly ascribing knowledge vary by social context). Fassio and McKenna's *primary* motivation for their revisionary proposal stems not from a concern about the efficacy with which the knowledge concept

fills the purposes for which it is used, but rather from the claim that there are a number of inconsistencies within the principles that govern 'knowledge' – a hypothesis that they adapt from Weiner (2009). For instance, the two example principles mentioned above are plausibly inconsistent, because the Social Principle allows that it might be correct to make differing knowledge ascriptions to subjects with identical evidence, so long as the social context differs – thus violating Parity of Evidence. Fassio and McKenna go on to argue, quite plausibly, that an inconsistent concept is one that ought to be revised.

Fassio and McKenna constrain their own revisionary tweaks by looking not only to how well knowledge's purposes are fulfilled, but also to retaining, as much as possible, the aforementioned principles. Their ultimate proposal is a concept, which they dub knowledge*, which is (roughly) a revisionary version of familiar subject-sensitive invariantist views. They also consider a revisionary *insensitive* invariantist view, but argue that it does not fulfill knowledge's functions as well as knowledge*. They argue, that is, that 'The best revisionary account of knowledge* is the account that is best suited to fulfilling the functions of knowledge' (Fassio and McKenna 2015, 15). They then claim that if two proposals fill the relevant functions equally well, then we should opt for the one which is most conservative with regard to the principles which characterize our current knowledge-concept.

Given that Fassio and McKenna hold our current knowledge-concept to serve a number of distinct purposes, however, this overall approach should strike us as a bit puzzling. Why must we limit ourselves to a single 'best' revisionary proposal? Why should we not introduce *multiple* successor concepts, each tailor-made to optimally fulfill some subset of the purposes that knowledge serves? Arguably, the very inconsistencies that Fassio and McKenna identify suggest that our current concept pushes in multiple directions. This may be a direct result of trying to do multiple jobs that call for conflicting characteristics – for a feature that enhances the knowledge-concept's suitability to fulfill one purpose (say, identifying reliable informants) may well decrease its suitability to fulfill another purpose (say, serving as the norm of action).

But the possibility of a pluralist strategy is, surprisingly, never mentioned by Fassio and McKenna as a potential option. I suspect that, in line with many advocates of revisionary approaches to philosophy, they hold that any legitimate successor to the concept *knowledge* must preserve most of its semantic content. And a successor concept tailored only to, say, *one* of knowledge's 'jobs' might well substantially differ in

meaning/content from its predecessor. The sort of similarity desideratum that Fassio and McKenna's monist approach suggests is of course present in Carnap, to whom Fassio and McKenna explicitly appeal. Fassio and McKenna may additionally be thinking here of the sort of Strawsonian concerns we discussed briefly above – they may think, with Strawson, that a revision of our knowledge concept must still be the concept *knowledge*, and not some new concept entirely.

On the approach to conceptual engineering I advocate, however, semantic similarity is not taken to be of any particular value outside of its tendency to reflect similarity of purpose. Nor is Strawson's problem held to be a serious challenge to the engineer. Insofar as Strawson's challenge has shown us that there ought to be *some* connection between the original concept and its engineered successors, I would argue that sufficient connection is provided by the purposes they share. Just as I am unmoved by concerns about changing the subject, I am not troubled at all by the idea that pluralist engineering *replaces* rather than *revises* its targets. Indeed, that's plausibly an *inevitable* consequence of one-to-many engineering – when an engineering intervention leaves us with multiple concepts, at most one will be 'the same concept' as the original. At most one, that is, will be a revision rather than a (partial) replacement.

Let's turn now to the roles our current concept of knowledge plays. Fassio and McKenna discuss three core purposes that the concept of knowledge plausibly serves: identifying reliable informants, licensing an inquirer to close inquiry, and serving as a norm for acceptable assertion and action. Let's briefly look at each of these roles in turn.

The first of these roles is inspired by the work of Edward Craig, who argues in his 1990 work *Knowledge and the State of Nature* that the primary purpose of the knowledge-concept is to enable identification of reliable informants – that is, good sources of testimony. Similar views have more recently been defended by Reynolds (2002) and by Hannon (2013, 2019). Identification of informants certainly does seem like a much-needed role for any human society - after all, a great many of our beliefs come from the testimony of others, and as such, it is crucial to be able to evaluate the accuracy of these testifiers and share that information with other potential inquirers. Reflection on ordinary uses of 'know' supports the contention that the concept of knowledge is frequently used for this purpose. As an example, when I ask my colleagues 'does anyone know what time the meeting starts?', it seems clear that my primary interest in deploying that concept is in identifying a reliable source of testimony. I'm not, for instance, obviously interested in

determining where to assign epistemic praise, or in cataloguing my colleagues' mental states. Similarly, for uses like 'you should talk to Loletta, she knows the department policies on this issue' or 'don't bother asking the cashier, he won't know whether this brand of cereal is gluten-free'. This all suggests that *one* central role of the concept of knowledge is to 'flag approved sources of information' (Craig 1990, 11).

The second of the roles that Fassio and McKenna discuss is that of serving as an 'inquiry-stopper'. As a brief gloss, the idea is that ascribing knowledge of some proposition P to a subject signals that the subject may appropriately close inquiry as to whether P. This suggested role for knowledge has been recently explored by Kappel (2010), Kelp (2011), and Rysiew (2012), all of whom present it as a potential improvement on Craig's approach – though from our current perspective, there's no reason why we shouldn't take both to be equally legitimate uses for the knowledge-concept. Again, it's easy to see why human societies would benefit from a concept that plays this inquiry-stopping role. Given that it's arguable that we do not have control over our belief formation³ – I cannot simply *will* myself to believe that $2 + 2 = 5$, even under Orwellian duress – the cessation of inquiry seems like a natural candidate for the primary target of epistemic praise and blame. Though I cannot control my beliefs directly, I *can* control whether I continue to pursue inquiry behaviors such as perceptual investigation, seeking out testimony, and the like. Thus, I may be legitimately blamed if I indulge in 'epistemic laziness' and end my inquiries too soon. Similar considerations lend a plausible case to the idea that inquiry-stopping may be the primary target of epistemic deliberation; I cannot decide whether to believe, but I can certainly decide whether to search for further evidence. As with the testimony role, everyday usage of 'knows' backs up the idea that our knowledge-concept is commonly used to signal whether inquiry may be closed. Imagine my partner asks me to look up our flight departure time. It's perfectly natural for me to reply, 'I don't need to, I already know it leaves at 5pm'. Similarly, it's perfectly natural for me to announce to a dinner guest, while cooking, something like 'I've just realized I don't know whether we have any sugar – I'd better check'.

The third (and arguably fourth) of Fassio and McKenna's roles for knowledge is serving as the norm of assertion and action. Giving again a rough gloss, to claim that knowledge is the norm of assertion is to

³For a classic entry point into debates over doxastic voluntarism, see Williams (1973).

claim that one may assert p only if one knows that p ; to claim that knowledge is the norm of action is to claim that one may use p as a premise in practical reasoning only if one knows that p .⁴ Both norms have numerous and varied defenders.⁵ As with the inquiry-stopper role, it's easy to postulate reasons why we might have a concept in place to serve these roles. Having a norm in place for assertion allows us to infer the degree of epistemic support the asserter holds – and, consequently, roughly indicates how much credence we should place in the asserted proposition. Having a norm in place for action is useful in guiding one's own practical deliberation, and also provides a means for holding others accountable for acts which, due to epistemic negligence, have caused undesirable consequences. And once again, reflection on everyday usage of 'knows' backs up the claim that knowledge serves these roles. Indeed, proponents of these norms – especially the knowledge norm of assertion – frequently point to evidence from linguistic usage as a core argument for supposing that knowledge plays such a role. For instance, Williamson (2000) notes that 'how do you know?' is an extremely common way to call the appropriateness of another's assertion into question; similarly, 'you don't know that' is commonly used to criticize unwarranted assertions. It is equally common to criticize *actions* by pointing to an agent's lack of knowledge, as in 'you shouldn't have pet that stray dog, you didn't know he wouldn't bite you'.

The purposes just surveyed are those which Fassio and McKenna focus on in their own revisionary project, but there are a few more nooks and crannies of the epistemological literature in which we might search for candidate uses for the knowledge concept. In particular, I would note Timothy Williamson's (2000) suggestion that knowledge is an irreducible component in successful psychological explanation and prediction. We will examine Williamson's arguments for this conclusion in a later section, but for now we can simply observe that certain uses of 'know' strongly suggest that we do currently employ the knowledge-concept for predicting and explaining the actions and mental states of our fellows. For instance, when I ask 'Does Jane know that we're meeting at the bus stop?', my primary interest seems to be in predicting her future behavior – as opposed to, for instance, determining whether she deserves

⁴I have phrased these norms as placing a *necessary* condition on proper assertion and action, rather than a sufficient or necessary-and-sufficient condition. The difference won't matter for our purposes, but see Brown (2008, 2010) for discussion of the alternative formulations.

⁵To give just a small selection, see Williamson (1996, 2000), DeRose (2002), and Turri (2010) for the norm of assertion and Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), and Hawthorne and Stanley (2008) for the norm of action.

epistemic praise, or whether she is a good source of testimony about our meeting point. I'm asking, rather, whether I should expect that she will indeed show up. Similarly, a natural response to a question like 'why did Mei step out to call her husband?' would be 'she knows he always worries when she gets home late'.

For all of the roles just surveyed, I've offered example natural-language uses that motivate the claim that we *currently* use the knowledge-concept for these purposes. But it should be noted that such linguistic evidence doesn't provide much support the further claim that we *ought* to use the knowledge-concept in this way. The fact that 'how do you know?' is the standard response to a questionable assertion is good evidence that knowledge *is* the norm of assertion, but it holds no weight against the engineer's suggestion that we might implement a *better* norm – one governed by a specially-designed epistemic concept. Similarly, the intuitive infelicity of sentences like 'P, but I don't know that P' provides descriptive evidence about our current assertion practices, but not prescriptive evidence about what norms we *ought* to endorse. *Mutatis mutandis* for arguments about whether or not the norm of assertion is constitutive of assertion, as claimed by Williamson (2000). From an engineering perspective, this is irrelevant – even if the knowledge norm *is* constitutive of assertion, the engineer can simply argue that we would be better off if our primary communicative speech act was assertion* instead. The majority of the epistemological literature on the roles just surveyed is primarily descriptive in the above sense, and is thus largely orthogonal to our current project; though as we'll see in the next section, some of the descriptive work can be adapted to motivate prescriptive claims.⁶

Let's sum up. We have, at current, identified a number of roles that our current knowledge-concept very plausibly serves: to flag approved sources of testimony, to signal that one may close inquiry, to regulate permissible assertion, to regulate permissible action, and to predict and explain the behavior and mental states of our fellows. This list is likely nowhere close to exhaustive, but I suspect we have more than enough on our plates for the moment. The roles we've surveyed here are enough to make the case that knowledge's current workload pulls it in potentially incompatible directions, and that ideal concepts for fulfilling

⁶In particular, some of the arguments offered in support of the descriptive thesis that knowledge is the norm of assertion/action/etc. appeal to normative considerations; these considerations may often be redeployable in service of an engineer's prescriptive recommendations.

these various roles might differ from the knowledge-concept in substantial ways.

Here's the strategy for the remainder of the paper. In the process of designing optimal successor concepts for knowledge, we will face decision points regarding ways in which our new concepts might either duplicate or depart from the various purported features of knowledge. Should the ideal concept for regulating assertion be factive? Should the ideal concept for playing knowledge's predictive/explanatory role invoke only internalist forms of justification? Fully exploring such questions would undoubtedly take multiple books, so I'll limit myself to a preliminary examination of a single feature that I take to give a good case for pluralism: sensitivity (or lack thereof) to attributor or subject circumstances. Specifically, I'll argue that the various roles we've identified for knowledge would be best served by concepts that differ along this dimension. For instance, I'll argue that while a concept that best serves as the norm of action would be sensitive to subject stakes, a concept that best fills knowledge's explanatory role would likely be invariantist. Towards the end of the paper, I'll also provide a few brief suggestions as to other features that might similarly provide fodder for the pluralist, such as factivity and the exclusion of epistemic luck.

3. Stakes-sensitivity and the roles of knowledge

There is, I think it's fair to say, no consensus whatsoever on whether our current knowledge concept is sensitive to features of the attributor's or subject's circumstances. The classic, traditional stance on knowledge is 'insensitive invariantism'; on this view, nothing in the attributor's or subject's circumstances is relevant to the truth of a knowledge attribution other than standard epistemic features such as whether the subject believes, has sufficient evidence, and so forth. Contextualist views, by contrast, hypothesize that the truth of a knowledge attribution is sensitive to features of the context in which it is uttered. On standard contextualist views, the truth-conditions for an utterance of the form 'S knows that p' become more demanding if (e.g.) the possibility of error is salient in the conversational context of the attributor. Such views are, therefore, 'attributor-sensitive'. Finally, a third type of view holds knowledge to be 'subject-sensitive'. On this view, the threshold for knowing is sensitive to non-epistemic features of the circumstances of the purported subject of knowledge, typically their practical interests. If the subject is in a high-stakes situation, their epistemic position must be quite high

before they count as knowing. But if they are in a low-stakes situation, the threshold for knowing is correspondingly relaxed. Such views are typically also labeled ‘invariantist’, in that they hold that the truth-conditions for ‘knows’ do not vary across *conversational* contexts.

Members of all three camps have argued in favor of many of the knowledge-roles discussed in the previous section. For example, there are insensitive invariantist (Williamson 2000), contextualist (DeRose 2002) and subject-sensitive invariantist (Hawthorne 2004) proponents of the claim that knowledge is the norm of assertion. As noted earlier, however, the dialectic in these debates is frequently based on ‘descriptive’ considerations such as linguistic usage, rather than prescriptive considerations. As conceptual engineers, we’ll approach from the prescriptive angle. Would an ideal concept for regulating assertion be insensitive, subject-sensitive, or attributor-sensitive? How about an ideal concept for regulating action? For licensing the closure of inquiry?

3.1. Stakes-sensitivity and the norm of action

Let’s start with action, for I think many of us find it highly intuitive that the appropriateness of acting on *p* is sensitive to the actor’s (that is, the subject’s) practical interests. If it does not much matter when my paycheck is deposited, it is fine for me to base my decision to put off a bank trip until Saturday on the basis of recollecting a Saturday bank-visit from several months prior. But if I face eviction if my check fails to clear on time, I really *ought* to double-check my bank’s hours before deciding to skip the bank on the way home. Note that neither of the aforementioned claims say anything about *knowledge* – they are instead common-sense claims about the degree of epistemic support that is required for rational practical decision-making.

It is *intuitive* that subject stakes factor into rational practical reasoning, but we can also provide argumentation in support of the reasonableness of this link. There is, after all, a trade-off to demanding higher epistemic standards in practical reasoning. Typically, attaining higher states of epistemic support requires an investment of time, cognitive effort, and even potentially material wealth (if one needs to, say, purchase a newspaper or run an additional medical test). When the stakes are low, this additional investment may not be worthwhile; but when they are high, it often is. A person suffering from obsessive-compulsive disorder may refuse to leave the house before they’ve checked multiple times that each appliance has been turned off; their epistemic position is improved by the

multiple-checking, but the corresponding expenditure of time and energy is excessive, given the circumstances. On the other hand, enforcing multiple checks on critical aircraft components prior to take-off and at intervals during a flight is worth the effort, given the consequences of failure.

Prima facie, then, an ideal concept for playing the ‘norm of action’ role – let’s call it ‘knowledge_{ACT}’ – would be subject-sensitive. Would it also be attributor-sensitive? It’s not clear why this would be a benefit. Given a subject-sensitive but attributor-insensitive concept, we’re provided with clear guidance towards determining whether an actor deserves epistemic praise or blame for their decision – they are blameworthy only if⁷ they fail to properly meet the epistemic standard set by their practical circumstances. But if knowledge_{ACT} were to incorporate attributor-sensitivity, the situation becomes less clear. Suppose that low-stakes Larry says ‘I know_{ACT} that p’. Suppose that his boss, high-stakes Harry, simultaneously says ‘Larry does not know know_{ACT} that p’. And finally, suppose that know_{ACT} incorporates attributor-sensitivity in such a way that both assertions turn out true. Now, when Larry uses p as a premise in practical reasoning, has he violated a norm? Is he deserving of epistemic blame? Making norm-violation (and consequently, epistemic blameworthiness) dependent upon the context of an attributor utterance leads us to say that the answer to this question depends on who’s talking. That seems to deny that there’s an ‘objective’ fact about the matter, rather in the manner of moral relativist views. At best it’s ‘true for Larry’ that he’s not blameworthy.⁸ This strikes me as undesirable – especially if epistemic blame has practical consequences, such as licensing later punitive attitudes or actions towards the subject.⁹ ‘Larry deserves to be fired for his poor judgment’ might end up false in Larry’s mouth, but true in Harry’s. God help the poor human resources employee assigned to adjudicate his wrongful termination appeal.

Should we instead ‘decouple’ knowledge_{ACT} attribution from assessment of proper action, so that Hank can truly say ‘Larry doesn’t

⁷Proponents of knowledge norms typically distinguish between norm-violation and culpability, arguing that some violations of knowledge norms may be blameless; the ‘only if’ phrasing here allows this by treating failure to have knowledge_{ACT} as a necessary but not sufficient condition to be blamed for a violation of the norm of action.

⁸Compare with subject-sensitivity. If knowledge_{ACT} is sensitive only to subject circumstances, all parties should agree about whether or not Larry is blameworthy – there is an objective right answer.

⁹Kauppinen (2018) makes very plausible suggestion that ‘epistemic punishment’ might involve e.g. lowering our willingness to believe future testimony from the subject and/or our confidence in their ability to properly conduct future inquiry or rational decision making.

know_{ACT} that p , but it was reasonable for him to act on p ? On such a view, the truth of a knowledge attribution would be attributor-sensitive, but the truth of an attribution of proper action would be invariant. We could thus incorporate attributor-sensitivity into knowledge_{ACT} without the undesirable implications just mentioned for blame-attribution. In fact, Hawthorne (2004, 85–89) suggests the corresponding view for ‘vanilla’ knowledge as the right interpretation of contextualism on cases like the one we’ve just discussed. He takes it as intuitively obvious that the normative facts are not attributor-sensitive, and thus holds that, on contextualist views, the sentence ‘you should rely on propositions that you don’t know to be true in your practical reasoning’ could in some contexts turn out true (Hawthorne 2004, 88). Hawthorne then notes, as an argument against the contextualist, that this sentence has a Moore-ish unintuitiveness.

As conceptual engineers, the linguistic infelicity of such sentences is for us beside the point. But what *is* to the point: if knowledge_{ACT}’s primary purpose is the assessment of rational action, attributor-sensitivity combined with ‘decoupling’ of this sort seems at best pointless, and at worst downright antithetical to knowledge_{ACT}’s suitability for its role. Nothing seems to be gained by adding in attributor-sensitivity other than, perhaps, an ability to capture certain contextualist intuitions about ‘knows’ – but capturing intuitions about ‘knows’ is not a desideratum of practical conceptual engineering. Worse, adding in attributor-sensitivity for knowledge_{ACT} attribution but not assessment of proper action seems to unduly complicate the concept, reducing user-friendliness – why use a more complex tool for the job at hand when a simple-to-use tool would do the same work just as well? I’d suggest, then, that attributor-sensitivity adds no benefit to an engineered action-regulating epistemic concept.

3.2. Stakes-sensitivity and the norm of assertion

Let’s turn now to assertion. Would an ideal norm of assertion require a level of epistemic support that is subject-sensitive, attributor-sensitive, or neither? I think a good case can be made that assertion should be stakes-sensitive, but that this sensitivity should be calibrated to the stakes of the *audience*, rather than the asserter. Consider the following imagined case from Jessica Brown:

[S]uppose that Sally is standing on the train platform waiting for a train to Edinburgh. Both express and non-express trains frequently pass the platform. She

has consulted the timetable and so comes to truly believe that the next train is an express to Edinburgh. She has a mild preference for taking an express train, but it will not be a disaster if she happens to take a non-express. This would merely result in her getting home 10 minutes later than need be. In such circumstances, it seems that she is in a good enough epistemic position to believe that the next train is an express and act on this assumption, say by taking the next train. Despite this, it need not be the case that she is in a good enough epistemic position to assert that the next train is an express. Imagine that a stranger approaches her on the platform and explains that it is crucial to him whether the next train is an express or not. Given how high the stakes are for the stranger, it seems that Sally is not in a good enough epistemic position to assert that the next train is an express. This is so even though it continues to be the case that, given her own much lower stakes, she is in a good enough epistemic position to believe that the next train is an express and act on it herself. Indeed, she may well tell the stranger that he should check at the very same time as she steps on to the train herself. (Brown 2012, 140)

My intuitions mirror Brown's on this case; but more importantly, reflection backs up this verdict. Assertion is a means for communicating information. When I make a (sincere) assertion that p , I typically aim for my audience to come to believe that p . And when I aim for my audience to come to believe that p , I typically do so because I expect them to use that information in future reasoning, especially practical reasoning.¹⁰ The sheer joy of acquiring new beliefs, though a lovely side benefit, isn't really the main purpose of assertoric communication. When I shout 'there's a car coming!' to a child playing in the street, my primary goal isn't to add to her store of beliefs – it's to get her to *move*.

The audience of an assertion will (again, typically) be aware of all this. So the stranger at the train station, assuming he takes Sally to be sincere and well-intentioned, will conclude from her assertion that she would recommend that he act on the proposition that the next train is an express. And Sally, if she is sincere and well-intentioned, should only assert that the next train is an express if she thinks the stranger ought to act on that claim. So she should only assert 'the next train is an express' if her epistemic position is high enough to license acting on that claim, given the *stranger's* stakes.

Let 'knowledge_{AST}' be our custom-built concept for playing the assertion-governing role. As a 'first draft' proposal, then, knowledge_{AST} should be sensitive in some way to audience stakes, such that Sally does not

¹⁰These two claims are not to be read as any kind of definition of assertion, or as proposing an individuating feature of assertion. It's enough for current argumentative purposes that *most* assertions are accompanied by such intentions, such that a hearer is justified in inferring said intentions (provided the hearer has no evidence of insincerity, etc).

know_{AST} that the next train is an express and is thus not in a position to make a norm-complying assertion to that effect to the stranger. She may, however, know_{ACT} that the train is an express. Note that knowledge_{AST}'s sensitivity is different from the standard sort of attributor-sensitivity postulated by contextualism. If a third party is watching the interaction between Sally and the stranger, the truth of 'Sally knows_{AST} that the next train is an express' as uttered by that third party is, on the current proposal, not sensitive to the third-party speaker's stakes.¹¹

What if the asserter has no evidence regarding the stakes of her audience, or has a justified false belief about those stakes? Does she nonetheless violate the norm of assertion if she fails to calibrate her assertion behavior to those unknown stakes? Should we instead hold that assertion should be governed by the speaker's *justified beliefs* about her audience's stakes? There are arguments to be made for either answer. On the one hand, it seems unfair to hold the asserter to a norm that she may not be in a position to follow, due to lack of evidence about her audience. But on the other, we can appeal to the distinction between violating a norm and being *blameworthy* for that violation – if an assertion is made on the basis of false but justified beliefs about the audience's stakes, this might be a blameless violation of the norm of assertion. Such a violation would be analogous to, for instance, an inexperienced traveler unintentionally violating an obscure local etiquette norm.

Moreover, there are arguments to be made in favor of an 'objective' norm rather than a 'psychologized' norm. By a 'psychologized' norm, I mean one for which norm-compliance is substantially determined by the character of the asserter's beliefs – in other words, a norm of the form 'assert p only if your evidence for p meets the epistemic threshold set by *what you justifiedly believe to be your audience's stakes*' rather than 'assert p only if your evidence for p meets the epistemic threshold set by your audience's stakes'. Most of our norms and laws seem to be 'objective' rather than 'psychologized' in this sense. Signs on public transport ask us to give up our seats for the pregnant and elderly, not for those we *justifiedly believe* to be pregnant or elderly. The law states that drivers should only enter the intersection when the light is green, not when they

¹¹A full proposal would need to give a definition of what it is to be the 'audience' of an assertion, but for current purposes a rough first stab will do: the audience consists of the persons such that the speaker intends them to hear (or read, or what have you) the assertion. Sally therefore is under no obligation to calibrate her assertions to a third-party eavesdropper's stakes. The question of how the threshold for knowledge_{AST} is set when the audience consists of multiple persons with differing stakes might be answered in different ways, but a plausible proposal is that the highest-stakes member of the group sets the threshold.

justifiedly believe the light is green. The law doesn't specify that it's ok to possess narcotics as long as you *justifiedly believe* that they are not prohibited. And so on.

Why should it be the case that we tend to utilize 'objective' norms? For one, objective norms are standardly simpler than psychologized norms – and are therefore easier to understand and to communicate to others. Relatedly, violations of objective norms are easier to detect than violations of psychologized norms, which require extra insight into the purported violator's mental states – it's easy to tell that a player is out-of-bounds in a ball game, but much harder to tell whether the player justifiedly believes they are out-of-bounds. Another, perhaps more weighty consideration in favor of objective norms comes from the motivating force of unintended norm-violation. If an asserter later comes to learn she was (justifiedly) mistaken about her audience's stakes, it would be preferable for her to make an attempt to retract or qualify her assertion, as in 'Oh, I didn't realize it was so important for you, perhaps you shouldn't take my word for it'. Deeming her previous assertion a (blameless) norm-violation motivates such behavior, for we are typically motivated to correct our mistakes – even when blameless. Deeming the assertion norm-conforming does not similarly motivate, for on such a view there is no mistake to correct.

The above considerations, if correct, would not only explain but also validate our preference for objective norms. Nonetheless, the arguments presented here are certainly far from decisive – it's possible that there are considerations that favor use of psychologized norms either generally, or in epistemology specifically. Thus, though I'm tempted to plump for an objective norm, I'll here leave the question open for future debate.

3.3. Stakes-sensitivity for testimony and for inquiry-stopping

Let's move now to testimony. Assertion is of course intimately linked to testimony, and so it's a real possibility that knowledge_{AST} might also be ideally suited to play knowledge's testimony-approving role. But we should consider possible complications. Hannon (2013) argues, for instance, in favor of a modified contextualism as best fitting knowledge's role in flagging reliable testifiers. On Hannon's view, the threshold for knowing shifts according to conversationally salient facts about the stakes of the person to whom a testifier is being recommended. But in the absence of conversationally salient information about said stakes, 'the standard for counting as a reliable informant will be set to a point

where one must satisfy the needs or interests of almost any inquirer' (Hannon 2013, 909). A 'default' standard makes substantial sense from a prescriptive perspective, for in flagging others as potential sources of information a chain of transmission will frequently be involved. Suppose I learn from Amit that Bronwyn knows a lot about the quality of our city's local doctors, and suppose that Amit learned this from Chris, who learned it from Diana. If Diana had merely transferred to Chris the information that Bronwyn has sufficient epistemic standing to meet *Chris's* testimonial needs, Chris may not have been in a position to recommend Bronwyn's testimony to Amit (who may have higher stakes). And similarly for Amit's recommendation to me. By contrast, if there is a general, agreed-upon 'default' threshold for reliable testimony, chains of testimony-recommendation avoid such issues.

Does this motivate introducing an additional concept, $\text{knowledge}_{\text{TES}}$, that is used to flag reliable testifiers? Or perhaps it simply motivates including a similar 'default' in $\text{knowledge}_{\text{AST}}$ – after all, when I acquire a new belief from another's assertion, I may later pass it on to another. Alternately, perhaps the optimal engineering solution is to design two concepts to jointly play a combined 'testimony-and-assertion' role, one sensitive to audience stakes and another serving as an insensitive 'default'. On such a proposal, the sensitive concept would be used in all testimony-assessment and assertion-assessment contexts when stakes information is available, and the invariant concept would be used in those contexts when such information is absent.

But then again, perhaps the assertion role and the testimony should be kept separate, with one concept used in assertion-assessment and another in testimony-assessment – for it's plausible that there may be cases where being a good source of testimony and following the norm of assertion come apart. Imagine, for instance, that high-stakes Harry overhears low-stakes Larry say to low-stakes Lisa that the next train is an express. Harry might speak truly when he says to himself 'Larry doesn't know $_{\text{TES}}$, I should double-check' – that is, Harry might correctly hold that Larry is not a good source of testimony (for Harry). And yet Harry might at the same time correctly hold that Larry was perfectly warranted in asserting as he did (and thus, that Larry knew $_{\text{AST}}$). The terrain here is tricky, and we unfortunately haven't got space to fully navigate it here. But at the very least, the above intricacies reflect just how complex knowledge's 'job description' really is. That alone suggests that we might benefit from implementing a little division of labor.

Similar potential overlap arises for knowledge's role as an 'inquiry-stopper'. I briefly noted in an earlier section that inquiry-stopping seems a natural locus for epistemic praise and blame, and for personal epistemic deliberation. This is because inquiry behavior, unlike belief formation *per se*, is something over which we have intentional control. However, epistemic praise and blame itself seems to be very frequently elicited by inappropriate action or assertion, suggesting that the inquiry-stopping role might be well-filled by concepts we've already postulated. If I aim to use p in practical reasoning, for instance, then I should stop my inquiry only when I possess knowledge_{ACT}. *Mutatis mutandis*, it seems, for asserting p and knowledge_{AST}. But what if I am simply curious about p , and have no immediate plans to act on or assert p ? In such cases, the relevant question seems to be whether my inquiry suffices to make it appropriate for me to *believe* p . We might, then, need an additional concept to flag whether my inquiry has been sufficient to license 'bare' belief which is not yet intended for use in assertion or practical reasoning. Would such a concept be sensitive to attributor or subject stakes? It's not immediately clear why it would be. If I'm not going to be acting on p , then my current stakes (or anyone else's) seem irrelevant to whether or not I should believe p ; similar for assertion. So perhaps here there is call for an invariant, insensitive concept – knowledge_{BEL}.

3.4. Stakes-sensitivity, prediction, and explanation

Let's finally look at knowledge's role in prediction and explanation. Williamson (2000) argues that knowledge plays an ineliminable role in explanation, one linked to the greater *stability* that knowledge enjoys when compared to other epistemic states. Williamson offers an imagined case in which a burglar spends hours ransacking a house, running a significant risk of being caught. The burglar's behavior can be explained if we suppose the burglar knows that there is a diamond hidden in the house. By contrast, the mere information that he believes truly is an inferior explanation of the burglar's behavior, because mere true belief may be consistent with, for instance, a lucky guess – in which case any reasonable burglar would have given up the chase after an initial search.

[T]he probability of his ransacking the house all night, conditional on his having entered it believing truly that there was a diamond in it, is lower than the

probability of his ransacking it all night, conditional on his having entered it knowing that there was a diamond in it. (Williamson 2000, 61)

Thus, knowledge provides (*mutatis mutandis*) a better explanation of the burglar's actions than mere true belief. The same consideration makes knowledge, in this case, a better predictor of said actions.

The argument, then, is that in certain cases where an actor displays perseverance in the face of apparent counter-evidence, knowledge's stability allows it to provide a better explanation of behavior than oft-fragile true belief. Does knowledge provide a better explanation in such cases than *justified* true belief? Williamson claims it does, for justified true belief is consistent with Gettierization. For example, the burglar's justified true belief that there is a diamond in the house might be based on (say) the false belief that the diamond is under the bed. But a burglar whose belief 'there is a diamond hidden in the house' is based on the false belief 'there is a diamond under the bed' is likely to give up the search once he's failed to find the diamond under the bed – for once the belief that there is a diamond under the bed has been falsified, he no longer has reason to believe that there is a diamond in the house. So the likelihood of the burglar ransacking the house, conditional on his knowing, is higher than the likelihood of the burglar ransacking the house, conditional on his merely truly, justifiably believing.

Let's grant that this argument is correct, and shows that there is a unique explanatory role played by our knowledge-concept. As noted earlier, natural language backs up this claim, for we frequently invoke knowledge in our explanations of others. So what would an ideal concept filling this role – let's call it knowledge_{EXP} – look like? Would it be sensitive to subject or attributor stakes? A sensitivity to attributor stakes seems undesirable, for the behavior of our fellows is not typically sensitive to our *own* conversational stakes. To use Williamson's example again, the likelihood that the burglar persists in his search is unaffected by our conversations about his epistemic standing (unless, of course, the conversation is loud enough to alert him to the presence of witnesses).

What about subject sensitivity? Here things are more subtle. Typically, proponents of subject-sensitive views on knowledge take the relevant facts about the subject's circumstances to be *objective*, in that the relevant facts are the subject's *actual* stakes and not the subject's beliefs about her stakes. Consider 'ignorant' high stakes cases, such as the variant on the classic bank cases found in Stanley (2005):

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Looking at the lines, Hannah says to Sarah, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.' (Stanley 2005, 5)

Stanley's intuition, shared by many, is that the bank-goers fail to know that the bank will be open on Saturday. But if we take this case to show that sensitivity to subject stakes means sensitivity to the *objective* stakes of the subject, then it seems $\text{knowledge}_{\text{EXP}}$ should not incorporate such sensitivity. After all, the behavior of Hannah and Sarah is not affected by facts they fail to know about their own stakes. On the other hand, a sensitivity to *subjective* stakes – that is, to the subjects' *beliefs* about the importance of being correct – might well be a predictive/explanatory boon. A subject who is very wary of error will likely behave in a different fashion than a subject with identical evidence but a more relaxed attitude towards failure. So if $\text{knowledge}_{\text{EXP}}$ incorporates any stakes-sensitivity, it should arguably be sensitivity to *subjective* subject stakes. If we hold that to not really be sensitivity in the relevant sense (after all, it is not sensitivity to the subject's external *circumstances* but to her own *doxastic states*, and that latter is something insensitive invariantists already accept), then it looks like $\text{knowledge}_{\text{EXP}}$ ought to be *insensitive*.

We can go further. Given that it is the *stability* of the actor's mental state that seems to be the key explanatory feature, our proposal for an optimal concept to play knowledge's predictive/explanatory role might eschew mention of stakes entirely. It might simply be something like 'being confident enough in p to continue basing one's actions on it in the face of any apparent counter-evidence one is likely to encounter'. Such a concept would differ from knowledge not only in being non-factive, but also in being compatible with at least some forms of Gettierization – ones which are such that the source of the Gettierization is unlikely to be revealed to the actor. Consider for instance a variant on Williamson's case modeled on Lehrer's (1965) Nogot case, where the burglar's belief is based on overhearing false-but-convincing testimony from the house owner that she has hidden a diamond somewhere in the house – while, unbeknownst to either, a diamond belonging to the previous owner is in fact concealed beneath the floorboards. This burglar

does not know, but is ‘confident’ in the sense specified above – for he is not likely to encounter any evidence that will undermine his belief.

In fact, critics of Williamson have already noted that a non-factive yet stable epistemic state such as confident belief suitably explains the burglar’s behavior – see for instance Kaplan (2003) and Magnus and Cohen (2003). We can perhaps reply on Williamson’s behalf that, as a matter of *descriptive* fact, people much more frequently invoke knowledge than ‘confident belief’ in actual explanation and prediction. But again, our project is prescriptive. And, given that knowing is compatible with non-confidence – just think of cases of ‘gaslighting’, where a victim is manipulated into doubting what she clearly knows – it’s arguable that the probability of burglar-type behavior conditional on confident belief is higher than the probability of burglar-type behavior conditional on knowing. Knowledge_{EXP}, then, might end up being a concept something like ‘confident belief’.

4. Conclusion

Though I’ve devoted the bulk of this paper to sensitivity, there are a number of other dimensions on which custom-made knowledge-successors might diverge from their predecessor – as the final paragraph of the previous section indicates. We’ve noted that an ideal knowledge_{EXP} concept might be non-factive, and compatible with at least some cases of epistemic luck. Similar suggestions might be made for knowledge_{ACT} and knowledge_{AST} – we might here look for inspiration to e.g. Lackey’s (2007) non-factive ‘reasonable to believe’ norm for assertion, or Brown’s (2008) argument that the norm of action permits acting while Gettiered. These are offered by their authors as descriptive accounts of our current norms, but we can – and should – look to their arguments in service of our prescriptive designs. Consideration of other epistemic features might ultimately lead us to modify these concepts yet further, leading to a set of successor concepts that may well diverge quite substantially from their predecessor, and from one another. Off the top of my head, potential features to look at might include safety, sensitivity, internal vs. external varieties of justification, quantity/quality of justification required, luminosity, and closure. Likely this list barely scratches the surface.

The discussion I’ve offered in this paper has been both very preliminary and very limited in scope. As usual, my initial ambitions for this paper far outstripped my word limit – and thus, to my chagrin, I haven’t even treated a single knowledge-feature as thoroughly as I would have liked. In particular, it’s worth noting that proponents of the various positions

discussed above – contextualism, subject-sensitive invariantism, and insensitive invariantism – can devise (and in many cases have devised) elaborations on their views that may well accommodate the apparent tensions that result from the variety of roles knowledge plays in our cognitive lives. I certainly haven't done much (or anything, really) to show that such strategies fail, or that a pluralist account is decisively superior to suitably sophisticated monist alternatives. But the *approach* I've presented here – that of *independently* considering which forms of sensitivity best suit certain roles, leaving open the possibility that the result will be a plurality of concepts – has not, to my knowledge, been seriously considered as an option. So I'll rest content if I've at least managed to add another contender to the ring – a plucky, untested newcomer with an unconventional right hook (to stretch the boxing metaphor a bit!).

Ultimately, my primary goal in this paper is methodological; it is to explore how conceptual engineering methods actually play out when applied to existing philosophical fields. My hope is that I've offered some indication of what a pluralism-friendly, practical approach to conceptual engineering would look like 'on the ground'. I further hope that this preliminary effort indicates how a shift in methodology might provide a new perspective for approaching old problems. Rather than trying to piece together a single analysis to do justice to the plausibly incompatible intuitions that drive invariantist, contextualist, and subject-sensitive views of knowledge, we can take these intuitions as indications that our current concept might be overburdened with jobs that make opposing demands – the conceptual equivalent of a runway model moonlighting as a competitive eater. Perhaps, then, it's time to consider making some additional hires.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by Research Grants Council, University Grants Committee [grant number HKU 13603718].

ORCID

Jennifer Nado  <http://orcid.org/0000-0003-1680-2975>

References

- Brown, J. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Nous* 42 (2): 167–189.
- Brown, J. 2010. "Knowledge and Assertion." *Philosophy and Phenomenological Research* 81 (3): 549–566.
- Brown, J. 2012. "Assertion and Practical Reasoning: Common or Divergent Epistemic Standards?" *Philosophy and Phenomenological Research* 84 (1): 123–157.
- Craig, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- DeRose, K. 2002. "Assertion, Knowledge, and Context." *The Philosophical Review* 111 (2): 167–203.
- Fantl, Jeremy, and Matthew McGrath. 2002. "Evidence, Pragmatics, and Justification." *The Philosophical Review* 111 (1): 67–94.
- Fassio, D., and R. McKenna. 2015. "Revisionary Epistemology." *Inquiry* 58 (7–8): 755–779.
- Hannon, Michael. 2013. "The Practical Origins of Epistemic Contextualism." *Erkenntnis* 78: 899–919.
- Hannon, M. 2019. *What's the Point of Knowledge? A Function-First Epistemology*. Oxford: Oxford University Press.
- Haslanger, S. 2000. "Gender and Race: (What) are They? (What) Do We Want them to Be?" *Noûs* 34 (1): 31–55.
- Haslanger, S. 2020. "How Not to Change the Subject." In *Shifting Concepts: The Philosophy and Psychology of Conceptual Variation*, edited by T. Marques, and A. Wikforss, 235–259. Oxford: Oxford University Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., and J. Stanley. 2008. "Knowledge and Action." *The Journal of Philosophy* 105 (10): 571–590.
- Kaplan, M. 2003. "Who Cares What You Know?" *The Philosophical Quarterly* 53 (210): 105–116.
- Kappel, K. 2010. "On Saying That Someone Knows: Themes from Craig." In *Social Epistemology*, edited by A. Haddock, A. Millar, and D. Pritchard, 69–88. Oxford: Oxford University Press.
- Kauppinen, A. 2018. "Epistemic Norms and Epistemic Accountability." *Philosophers' Imprint* 18 (8): 1–16.
- Kelp, C. 2011. "What's the Point of 'Knowledge' Anyway?" *Episteme; Rivista Critica Di Storia Delle Scienze Mediche E Biologiche* 8 (1): 53–66.
- Lackey, J. 2007. "Norms of Assertion." *Nous* 41 (4): 594–626.
- Lehrer, K. 1965. "Knowledge, Truth and Evidence." *Analysis* 25 (5): 168–175.
- Magnus, P. D., and J. Cohen. 2003. "Williamson on Knowledge and Psychological Explanation." *Philosophical Studies* 116 (1): 37–52.
- Prinzing, M. 2018. "The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Objection." *Inquiry* 61 (8): 854–880.
- Reynolds, S. 2002. "Testimony, Knowledge, and Epistemic Goals." *Philosophical Studies* 110: 139–161.
- Rysiew, P. 2012. "Epistemic Scorekeeping." In *Knowledge Ascriptions*, edited by J. Brown, and M. Gerken, 270–293. Oxford: Oxford University Press.

- Scharp, K. 2013. *Replacing Truth*. Oxford: Oxford University Press.
- Simion, M., and C. Kelp. 2019. "Conceptual Innovation, Function First." *Noûs*, doi:10.1111/nous.12302.
- Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Strawson, P. F. 1963. "Carnap's Views on Conceptual Systems Versus Natural Languages in Analytic Philosophy." In *The Philosophy of Rudolf Carnap*, edited by P. A. Schilpp, 503–518. La Salle: Open Court.
- Thomasson, A. L. 2020. "A Pragmatic Method for Conceptual Ethics." In *Conceptual Ethics and Conceptual Engineering*, edited by H. Cappelen, D. Plunkett, and A. Burgess, 435–458. Oxford: Oxford University Press.
- Turri, J. 2010. "Epistemic Invariantism and Speech Act Contextualism." *Philosophical Review* 119 (1): 77–95.
- Weiner, M. 2009. "The (Mostly Harmless) Inconsistency of Knowledge Ascriptions." *Philosopher's Imprint* 9: 1–25.
- Williams, B. 1973. "Deciding to Believe." In B. Williams, *Problems of the Self*, 136–151. Cambridge: Cambridge University Press.
- Williamson, T. 1996. "Knowing and Asserting." *The Philosophical Review* 105 (4): 489–523.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.