


ARTICLE

# Wrongful ways to raise the epistemic standard

Jumbly Grindrod\* 

University of Reading, Reading, UK

\*Corresponding author. Email: [j.grindrod@reading.ac.uk](mailto:j.grindrod@reading.ac.uk)

(Received 17 December 2019; revised 29 June 2020; accepted 20 July 2020;  
first published online 1 October 2020)

## Abstract

This paper is concerned with identifying and accounting for cases where the epistemic standard is raised inappropriately. The first section is concerned with identifying a notion of a variable epistemic standard that is neutral regarding a range of theoretical issues. The second section argues that the possibility the epistemic standard could be raised in some epistemic inappropriate way warrants further investigation. The third section outlines and provides a partial explanation of such a case: one in which a climate change denier attempts to raise the epistemic standard in order to shut down inquiry.

**Keywords:** Epistemic standard; epistemic value; context-sensitivity; epistemic normativity

## 1. Contextual variability

The idea of an individual raising the epistemic standard will be familiar to many; the idea that one may raise the epistemic standard *wrongfully* less so. In order to broach the idea of wrongful raisings, we will need to clarify what is meant by ‘epistemic standard’ for the purposes of this paper. One problem in even discussing the idea of a variable epistemic standard is that it is nearly always tied to a particular theory that is controversial to some extent e.g. contextualism, pragmatic encroachment, semantic relativism etc. By contrast, my approach in this paper will be to begin with a notion of the epistemic standard that captures the idea that it can vary across contexts while remaining neutral on a number of theoretical issues. In doing so, the conclusions of this paper can remain neutral regarding, for example, the truth of epistemic contextualism. The remainder of this section will be devoted to developing this suitably neutral notion.

It is best to understand the notion of an epistemic standard by first considering the apparent contextual variability of knowledge-attributing sentences. *Context-shifting experiments*, such as DeRose’s bank case (DeRose 1992) and Cohen’s airport case (Cohen 1999) show that a sentence of the form “S knows p” appears true in one context and false in another. This is so even though the two contexts do not vary in the subject, the purportedly known proposition, and the epistemic standing the subject is in regarding that proposition.<sup>1</sup> Typically, the differences between the two contexts lie only in the practical stakes in being right or in the kind of alternatives to the proposition that have

<sup>1</sup>Here, “epistemic standing” is used as a kind of placeholder for the truth-relevant features that, together with truth and belief, determine whether a subject truly knows. The reader is free to replace this phrase with their favourite third condition, such as justification, warrant, evidence, reliability, safety, etc.

been made salient. This variation in the apparent truth value of knowledge attributions stands in need of explanation, as traditional invariantist views would assign the same truth value to each knowledge claim.

One way to explain this is to adopt epistemic contextualism – to claim that the truth-conditional content of knowledge attributions can vary across different contexts of utterance, and that the above cases are cases where such variation occurs. However, while many will be familiar with this line of reasoning towards contextualism, few would take such a line to settle the matter. The problem is that there are many other ways to explain the context-shifting experiments. Firstly, there are pragmatic encroachment views of knowledge, such as subject-sensitive invariantism (SSI). Whereas contextualism is primarily a semantic view on the truth conditions of knowledge attributions, pragmatic encroachment views are primarily metaphysical views on the nature of knowledge. For example, SSI claims that knowledge is a tertiary relation between a subject, a proposition, and the practical stakes the subject has in being right about the proposition (Hawthorne 2004; Stanley 2005; Fantl and McGrath 2009). This kind of view can capture the context-shifting experiments without thereby adopting a contextualist semantics, provided that there is some variation in the practical stakes of the subject.<sup>2</sup> Secondly, there are pragmatic approaches, which provide a pragmatic explanation of the intuitive difference in truth value across the context-shifting experiments. For example, it might be claimed that if Keith utters “I do know the bank will be open” in a high standards context, what he says is true (contrary to intuition), but he thereby flouts the Gricean maxim of relation (i.e. be relevant), and so communicates the related proposition that he can rule out certain alternatives that are salient in that context (Rysiew 2001, 2005). Such a view can capture the intuitive truth values of the knowledge attributions in each context while denying the contextualist claim that there is a context-sensitive semantics. Thirdly, there are epistemic relativist views that will claim that the truth values of knowledge attributions can vary across different *contexts of assessment*. For the relativist, it is not that the content communicated by uttering a knowledge sentence is dependent upon the context it is uttered in; the content remains invariant. Instead, when it comes to assessing the truth of such an utterance, the proper outcome of the assessment is dependent upon certain parameters that the assessment is made in. In one important sense, the relativist posits a kind of (truth value) context-sensitivity that outstrips the contextualist view. They will claim that there can still be a variation in truth value even if all features of the context of utterance remain fixed (relativists often illustrate this by appealing to *retraction* cases (MacFarlane 2011)). These three alternatives illustrate that there are numerous theoretical options on the table, and that as a result, arguing directly from context-shifting experiments to epistemic contextualism (or any one of these alternatives) is not plausible.<sup>3</sup>

Furthermore, which of these broad approaches is more plausible depends a great deal on the further detail that is added to them. For example, contextualist and relativist views are plausible to the extent that they provide a plausible semantics, and pragmatic approaches are plausible to the extent that they provide a plausible pragmatics. This

<sup>2</sup>Assuming that pragmatic encroachment views will adopt a semantics that reflects the underlying metaphysics (not a trivial assumption, but not one that is easy to reject either), then pragmatic encroachment views will adopt a semantics that allows for a variation in truth value across contexts where the practical stakes of the subject of the knowledge attribution varies (as in the Bank case and Airport case). The context-sensitivity predicted by contextualism is broader than this: contextualism will claim that there can be variation across contexts of utterance even when all relevant facts about the subject (including her practical stakes) remain fixed.

<sup>3</sup>See Gao *et al.* (2017) for further criticism of DeRose’s “methodology of the straightforward”.

makes philosophical discussion on these issues difficult not only because they partly rely on further issues in related empirical fields (e.g. linguistics and psychology), but also because they are affected by issues within the related fields that are of a much more general nature than the topic of knowledge attributions. For example, the plausibility of epistemic contextualism will depend on more general questions regarding context-sensitivity: how pervasive is context-sensitivity in natural language? And are there specific tests for identifying cases of context-sensitivity? Our final position on the plausibility of contextualism will not be settled until we can find answers to these more general questions, and a similar point applies to all of the alternative views.

Described in this way the state of play regarding such theories can seem quite desperate. However, I am in fact optimistic that such issues can be resolved in the future. For the purposes of this paper, we are interested not in what divides these views but what unites them. All of the above views allow that there is a kind of contextual variability in the way we attribute knowledge. Yet far less attention has been paid to how this fact alone could impact upon, or be impacted by, other areas of epistemological interest, such as the value problems of knowledge or epistemic harms and injustices.<sup>4</sup> One clear reason for this is that thus far we have only been able to frame the discussion of this epistemic variability in the terms of one of the above theories, each of which are controversial. Consequently, an investigation of the type I have described – focused as it is on wrongful raisings of the epistemic standard – appears uninteresting because it is based upon a premise that is not obviously acceptable.

But things need not be this way. The solution lies in finding neutral terms with which to describe the phenomenon in question, which will enable us to investigate the phenomenon in such a way that it does not matter whether contextualism, SSI, relativism, etc. is true. We need a way to describe the difference between the two contexts in cases such as DeRose's bank case and Cohen's airport case. Why, for example, does it seem true for Keith to say that he knows the bank will be open in the low context, but false to say that he knows in the high context? One approach would be to remain strictly committed to the merely descriptive claim that the intuitive truth values vary across the two contexts. Of course, this would still be subject to confirmation via observation of our intuitive truth values, and this is something that has already occurred to some extent in the experimental literature (Hansen and Chemla 2013; Grindrod *et al.* 2018). In this paper, I will adopt a slightly stronger claim: that not only does it *appear* that there is a variation in communicated content across the two cases, but that there *really is* a variation in communicated content. The claim I am interested in can be outlined as:

*Modest variability:* What a knowledge attribution communicates about the epistemic position that the subject is in can vary across contexts of utterance even while the following remain fixed: the subject, the proposition claimed to be known, the truth of the proposition, and the epistemic standing that the subject is in regarding the truth of that proposition.

Modest variability is an interesting claim about our epistemic practices, and it is one that all of the above views agree upon.<sup>5</sup> In the next section, I hope to show how modest

<sup>4</sup>A couple of notable exceptions are Fantl and McGrath (2009: Ch. 6) and Gerken (2019). See also Grindrod (2019).

<sup>5</sup>This does not mean it is entirely uncontroversial. For example, more recently a number of views have defended the claim that what is communicated does not vary across the two contexts, but that the intuitions that we hold regarding such cases are affected by certain general heuristics that can explain why our

variability can be theoretically useful insofar as it allows us to investigate further epistemological issues in a relatively neutral manner. In the remainder of this section, I will discuss certain features of modest variability further.

First, it should be noted that in seeking more neutral terms, I have resorted to a communicative claim about what knowledge attributions communicate. This is a claim that all of the previously discussed views would accept insofar as they attempt to explain the bank case, and the bank case and other context-shifting experiments are cases where the intuitions in question are focused on communicative acts. Of course, pragmatic encroachment views of knowledge do not primarily phrase their view in terms of communicated content, but if their view is going to provide an explanation of the bank case and other context-shifting experiments, it will be under the assumption that the acceptability of attributing knowledge is tracking the underlying metaphysical issue of whether the subject in question has knowledge. All of the other views I considered are views about what knowledge attributions communicate. So in using these terms, we can capture a truth that all of the above views accept.

Note also that the notion of “communicated content” remains neutral about whether the content in question is generated by semantic or pragmatic means. At this point an objection could be raised that modest variability is *too neutral* and that it is thus entirely uncontroversial, simply because what we communicate using *any* sentence will vary across contexts. Take a mathematical statement such as “ $2+2=4$ ”. Mathematical statements like this are usually thought to be the best example of a sentence that communicates invariant content. If there is a non-empty set of context-invariant terms in a language, then it seems plausible that mathematical terms would be included as members. Yet what is communicated by the above statement can nevertheless vary across contexts, and this is simply due to the ability of speakers to generate highly particularised conversational implicatures. Consider Jack and Jill and Mike and Molly, who are looking for a restaurant for a double-date. They go to a busy restaurant and ask the maitre d’ if they have a table available. The maitre d’ states that they have a table for three. Annoyed, Jack says “Well,  $2+2=4$ ”. In doing so, Jack is communicating that a table for three is no good and perhaps that the maitre d’ should not have mentioned the table for three. This suggests that *any* sentence can vary in its communicated content because for any sentence there are particular contexts in which speakers can use that sentence to communicate highly particularised messages. So, the objection goes, modest variability does not pick out something interesting about our epistemic practices; it ascribes an uninteresting property to knowledge sentences that is possessed by any sentence in a language.

It is important to see how modest variability avoids this objection. Modest variability states that what knowledge attributions communicate *about the epistemic position that the subject is in regarding the particular proposition* can vary across contexts. That is, it claims that the variation in content is not in terms of some issue entirely divorced from attributing knowledge, but the variability concerns the central topic that knowledge attributions are about. So of course, knowledge attributions can be used to communicate highly particularised implicatures, and in this sense it is uncontroversial that their communicative content can vary. But the interesting claim we are concerned with is that

---

intuitions are wrong. For example, Williamson (2005) argues that the bank case can be explained via appeal to the *availability heuristic*, which is that people tend to think that if a possibility is made salient, then it is more likely (e.g. plane crashes and shark attacks are often perceived to be more likely than they in fact are). See Nagel (2010) and Gerken (2013) for similar appeals to heuristics.

what knowledge attributions communicate about a given epistemic position can vary, regardless of whether this is generated semantically or pragmatically.<sup>6</sup>

There is good reason to think that the level of description employed by modest variability, while more neutral than the terminology standardly used in the existing literature, is in fact the appropriate level of description for what might be described as the traditional approach to epistemology. One classical project in epistemology is to undertake a conceptual analysis for KNOWS i.e. provide necessary and sufficient conditions for the application of the concept. It is often acknowledged that one form of evidence that philosophers draw upon in order to correctly capture this is our linguistic intuitions – roughly, when it would be appropriate to attribute knowledge and when it would be appropriate to deny it. But, as Chomsky and others have noted, the bare linguistic intuitions we possess are blind to substantive theoretical matters such as semantic and pragmatic content:

We may make an intuitive judgment that some linguistic expression is odd or deviant. But we cannot in general know, pretheoretically, whether this deviance is a matter of syntax, semantics, pragmatics, belief, memory limitations, style, etc., or even whether these are appropriate categories for the interpretation of the judgment in question. It is an obvious and uncontroversial fact that informant judgments and other data do not fall neatly into clear categories: syntactic, semantic, etc. (Chomsky 1977: 4)

Some already recognise this. For example, Pritchard (2012) provides a taxonomy of the intuitions that epistemologists draw upon in the project of providing an account of knowledge. For our purposes, he distinguishes between *extensional intuitions* – intuitions about particular cases regarding when knowledge is present and when it is absent – and *linguistic intuitions* – intuitions regarding when it is appropriate or inappropriate to attribute knowledge. One may worry that there isn't an important sense in which these two sets of intuitions come apart, but Pritchard argues that they do come apart in cases where attributing knowledge would also communicate some pragmatic content e.g. implicated content (Pritchard 2012: 95). So epistemologists do often appeal to our linguistic intuitions as one of the evidence bases to form our theory around, but as Pritchard and Chomsky recognise, this evidence basis is neutral to the more specific issues of semantics, pragmatics, and otherwise. As a result, claims like modest variability more faithfully capture the evidence base that philosophers draw upon.

Once we have accepted modest variability, it becomes natural to think of the variability in terms of a variable epistemic standard. Contextualists, for example, usually claim that the epistemic standard relevant to the truth of knowledge attributions can be higher

<sup>6</sup>An anonymous reviewer questions whether certain pragmatic accounts of the apparent variability of knowledge attributions would be committed to modest variability. For instance, a pragmatic account might claim that in high stakes contexts, a knowledge attribution would conversationally implicate that the subject can rely on the proposition in question given their practical purposes or their stakes. And this implicated proposition, it might be claimed, doesn't strictly communicate anything about the subject's epistemic standard, it is only a claim about what propositions the subject ought to rely on. Now as has been widely recognised in the pragmatic encroachment literature, whether the subject ought to rely on the truth of a proposition, given their practical purposes, depends on how strong the subject's epistemic position is with regard to the proposition. It is in this respect that I take there to be variation in what is communicated about the subject's epistemic position even on the above view. But I recognise here that ideally the phrase "communicates about the subject's epistemic position" would be replaced with one that more clearly picks out the kind of communicative variation that is of epistemic interest. As it stands, I take the phrase to work well enough for the purposes of the paper.

or lower across contexts, and it is this that differs across various contexts. We can capture modest variability in terms of a variable epistemic standard as follows: in attributing knowledge, we communicate that a subject's epistemic position meets the epistemic standard, and this standard can be more or less demanding across different contexts of utterance. Thus, the variability in what knowledge attributions communicate can be described in terms of a variable epistemic standard. It is this notion of the epistemic standard that will be used throughout the remainder of this paper. Employing this notion allows us to investigate the phenomenon in question – inappropriate raisings of the epistemic standard – in a way that is not premised upon the truth of some controversial epistemological theory.

In this section, I have outlined modest variability as a more neutral claim that captures a fact about our knowledge-attributing behaviour and that is widely accepted across the literature. This variability is naturally understood via appeal to an epistemic standard that varies across contexts. The key point of this section has been that in employing a less committed notion of the epistemic standard, we become freed up to explore its impact upon our epistemic lives in a way that is not premised upon a particular controversial theory. The remainder of this paper serves as an illustration of that point. I will identify and diagnose a kind of epistemically unacceptable behaviour related to the varying epistemic standard. To do so, it is useful to first consider the interaction between varying epistemic standards and a much-discussed sceptical argument.

## 2. The sceptic as bully

Once we allow that there is a variable epistemic standard that at least affects what we communicate in attributing knowledge, the question arises as to whether there is a sense in which this variable epistemic standard could be misused in some way. Are there certain normative principles surrounding the epistemic standard, or certain harms or even injustices that can arise in relation to the epistemic standard? It may be that, while the epistemic standard serves its function in a good deal of cases, it is nevertheless prone to manipulation by nefarious actors. Brister (2009, 2017) has argued that if we accept the contextualist account of sceptical arguments, the sceptic can be viewed as this kind of nefarious actor insofar as they attempt to raise the epistemic standard to an extreme level. Consider the much-discussed closure-based sceptical argument:

1. I don't know I'm not a brain in a vat.
2. If I know I have hands, then I know I'm not a brain in a vat.
3. Therefore, I don't know I have hands.

The argument appears valid and the two premises seem difficult to deny. Yet the conclusion is a form of scepticism that appeals to few. According to the standard contextualist explanation, premise 1 *raises* the epistemic standard to a level where we are now considering brain-in-a-vat scenarios, and so, because a brain-in-a-vat scenario is not something that we have any evidence against, 1 is true in the sceptical context. But this only shows that the argument is effective within the confines of the sceptical context. In more ordinary contexts, where a less demanding standard is in play, 3 would not be true. Certainly, this kind of explanation does raise further questions,<sup>7</sup> but for our purposes the important point is that according to this explanation, the sceptic forces a high epistemic standard upon us by appealing to remote possibilities that we cannot

<sup>7</sup>Two pertinent questions are: (i) is 1 true in less demanding contexts? (ii) Can we still hold on to a principle of epistemic closure on this account?



rule out. Brister notes that there is an element of hostility in this behaviour, that the sceptic's move of raising the epistemic standard to such a high level is an "arbitrary deployment of epistemic power" (Brister 2017: 64). Elsewhere, she states: "from this perspective skepticism is a kind of power grab that puts other epistemic agents on the defensive and at a disadvantage. While some might believe that knowledge is power, the skeptic recognises that *denying* knowledge is power" (Brister 2009: 682). Something like this idea will guide the remainder of this paper, but given what was said in the previous section, we need to view the sceptic through more neutral terms. We can agree with the contextualist account that the sceptic, in employing the above argument, seeks to raise the epistemic standard, and that this in turn affects what we communicate in attributing knowledge. But as per modest variability, we remain neutral regarding the truth-conditional content of 1–3. We are only interested in the claim that in raising the epistemic standard, the sceptic affects what sentences like 1–3 communicate, and that this can be viewed as aggressive, as a kind of power move.

A natural response to Brister's description of the sceptic is that this isn't exactly a bully we should be worried about, for two reasons. First, we rarely, if ever, meet sceptics in everyday life, and while they do seem to show up on occasion in philosophy seminars, I think it is reasonable to doubt whether they manage to consistently live by their professed scepticism outside of the classroom. So while we may have to tread carefully in philosophical corridors, the immediate danger of meeting this brand of bully is minimal and thus not particularly interesting as far as our theory of epistemic wrongdoings goes. Second, even if sceptics were two-a-penny, it doesn't seem like the power they wield is particularly great. All the sceptic achieves is to affect what we communicate by attributing knowledge within a context; they haven't affected our credence or evidence or obviously harmed us, nor have they affected our reputation as epistemic agents.<sup>8</sup> So we shouldn't worry about the sceptic because (i) they're not exactly a common occurrence and (ii) their power to raise the epistemic standard does not seem particularly significant.

In response to the first reason, it is worth bearing in mind that while philosophical sceptics may be quite rare, the behaviour of the sceptic may be an extreme instance of a kind of behaviour that we do find in more ordinary discourse. In the next section, I will attempt to outline particular cases that I take to resemble the sceptic's behaviour in important ways.

In response to the second reason, it is important to keep in mind that even if the sceptic's power – to affect what is communicated by attributing knowledge – is ultimately a linguistic one, it doesn't follow that it is unimportant. In recent years, there has been an increased focus on the idea that philosophy should not only keep track of what concepts we do employ, but what concepts we *ought to* employ (Burgess and Plunkett 2013a, 2013b; Cappelen 2018). Burgess and Plunkett label this broader project *conceptual ethics*. There are two motivations behind this project. One is more theoretical and has its roots in Carnap (1950); the idea is that to theorise precisely on a given topic, we need to employ theoretical concepts that are not incoherent, problematically vague, or problematically ambiguous. Some of our ordinary concepts may not meet these criteria, and so we must *explicate* such concepts for theoretical purposes. The other motivation is more practical, and has its roots in the work of Haslanger (2000), among others. The idea is that the concepts we employ affect the way we think about a given topic, and in so doing

<sup>8</sup>This is particularly pertinent given that, in line with modest variability, no particular claim is made here about whether the sceptic affects semantic or pragmatic content, or whether the sceptic robs us of our knowledge in some metaphysical sense. All that is claimed is that the sceptic affects what is communicated by knowledge statements.

will impact upon our decisions and actions on that topic. If that is true, then some concepts affect us in bad ways and some in good ways.<sup>9</sup> This issue is perhaps most clearly seen in forms of discourse where a great deal rides on using the correct concepts. The concept of “rape” is a good example: in many countries the legal definition of “rape” has been redefined so that it can apply to an occurrence between a married couple. Doing so has been important not only so that incidences of marital rape are treated within the law as instances of rape, but so that the wider perception and discourse surrounding marital rape could be changed. In this respect, it is clear that the revised definition is an improvement on the older one, and this illustrates clearly that there are ethical questions about the concepts we employ.<sup>10</sup>

This practical motivation can plausibly be generalised beyond cases like ‘rape’ – where the potential harm involved is so severe – so that the question of conceptual ethics can be asked for a very wide array of terms:

Arguably, our conceptual repertoire determines not only what beliefs we can have but also what hypotheses we can entertain, what desires we can form, what plans we can make on the basis of such mental states, and accordingly constrains what we can hope to accomplish in the world ... Conceptual choices and changes may be intrinsically interesting, but the clearest reason to care about them is just that their non-conceptual consequences are pervasive and profound. (Burgess and Plunkett 2013a: 1096–7)

This alone should make it unsurprising that the question of conceptual ethics could be applied to our knowledge discourse. That is, there is not only the question of what our knowledge attributions mean, but what they ought to mean.<sup>11</sup> This becomes all-the-more pressing if we allow that knowledge plays some important role in our epistemic lives and this is one claim for which we can find broad agreement in the literature. Some have argued that knowledge is the basis for assertion or action; some have argued that knowledge serves to terminate inquiry; some have argued that knowledge serves to tag credible informants, and some have argued that knowledge possesses a distinctively epistemic value beyond its use in practical reasoning. But what the above views all agree upon is that knowledge plays an important role in our epistemic lives, and once we acknowledge this importance, the project of trying to understand what knowledge is for and which concept would best serve that purpose seems well-motivated.

The question of conceptual ethics is usually viewed as a revisionary one insofar as we are considering the possibility of changing our otherwise stable conceptual apparatus. But when we are considering context-sensitive communication, as we are in the case

<sup>9</sup>That this form of conceptual engineering rests on the idea that concepts can affect us in *good* and *bad* ways reveals that this project requires some form of prior theory of value, of what is good and bad. In domains where the value theory will also be contested (such as in the moral domain) it becomes difficult to engage upon this project of engineering our concepts so that they affect us in good ways. But there may still be cases where it is uncontroversial that a concept could be improved. Thanks go to Brad Hooker for emphasising this point.

<sup>10</sup>On one view of what concepts are, the idea that we could modify concepts does not really make sense. If we hold the view that a concept is defined by its conditions of application, then it doesn't really make sense to ask whether a concept would be improved by altering those conditions of application. So clearly the conceptual ethics project must delineate concepts some other way. Perhaps the clearest alternative is to view concepts in terms of the functional role that they play. On this view, we ask of a concept C what role it plays in our cognitive lives, and then we consider how C could be modified so as to best fulfil that role.

<sup>11</sup>Indeed, the concept of knowledge is the subject of Haslanger's revisionist programme (Haslanger 1999).



of knowledge attributions, there is the related question of how we ought to set the meanings of our context-sensitive expressions for the purposes of a conversation. The question is not whether we should seek a linguistic change in the meaning of “know”, but how we should set the meaning of the term within a context. Plunkett and Sundell have argued that the question of conceptual ethics can arise just as easily for context-sensitive content as it can for invariant content. Considering the case of the context-sensitive term “spicy”, they write:

An agreement amongst all the cooks in the kitchen that the chili can be described as “spicy” plays an important role in collective decision-making. In particular, it plays an important role in decision-making about whether to add more spice. (Plunkett and Sundell 2013: 15)

In this way, setting the meaning of “spicy” within a context is going to be important because it will affect set boundaries in terms of permissible actions. If we set the standard of spice to be very high, then this will affect when the criticism of something being “too spicy” comes into play. Plunkett and Sundell subsequently argue that the importance of setting the context-sensitive meanings of terms can give rise to *metalinguistic disagreements* over how to use a term, and that the phenomenon of metalinguistic disagreements can be used to provide an alternative explanation of the faultless disagreement cases that have been elsewhere used to argue for semantic relativism about particular expressions (such as taste predicates). Similarly, Ludlow (2014) has argued that, because context-sensitivity is pervasive, metalinguistic disagreements over what a term ought to mean for the purposes of a context can arise for nearly any term in a language, and he includes examples of disagreements surrounding the meaning of “sandwich”, “planet”, “athlete”, and many more terms.

The important point for our purposes is not whether context-sensitivity is pervasive, nor whether knowledge attributions are context-sensitive in a way similar to “spicy” and other terms, nor whether metalinguistic disagreements provide a good explanation of faultless disagreement cases. The important point here is that once we allow that knowledge attributions vary in what they communicate, and we allow that knowledge discourse is of particular importance, it becomes clear that the question of what knowledge attributions communicate within a context is an important one. Knowledge discourse performs an important function in our epistemic lives, and so what knowledge attributions mean within a context will affect the extent to which knowledge discourse can play that role.

Returning to the actions of the sceptic, we now have reason to be more suspicious of their behaviour. If there is an importance in what our knowledge attributions communicate, and sceptics are changing what knowledge attributions communicate to fit their own purposes, then we can see now that the sceptic’s actions are not inconsequential but could actually have negative effects on our epistemic lives. And while the sceptic might be something of a rarity, as I argued earlier, the sceptic’s behaviour may be an extreme form of a more general behaviour that is much more common. *Attempts to raise the epistemic standard* may be an ordinary part of our epistemic lives and there may be right and wrong ways to do this. In the following section, I will focus on one such case and will attempt to provide a partial diagnosis of what has gone wrong.

### 3. Wrongful standard-raising

In order to motivate the idea that there can be wrongful raisings of the epistemic standard, I will take my cue from Brister’s discussion of the climate change denier:

The climate science community considers the denialist strategy to be specifically designed to side-line inquiry away from practical goals and to engage scientists in a quixotic conversation about standards of justification, such as what counts as scientific consensus and whether there is a scientific conspiracy. (Brister 2017: 65)

There is a clear kind of epistemic behaviour often displayed by the climate change denier where an attempt is made to claim that the evidence provided in favour of climate change does not meet the standard that they think we ought to operate by. The following case outlines the kind of behaviour I have in mind:

**Climate change:** John is a climate skeptic and Sally is not. Sally says “Of course we can know that man-made climate change is occurring. The heat-trapping effect of CO<sub>2</sub> is well-established. The CO<sub>2</sub> levels are the highest they have been for millions of years, and we are starting to see the predicted effects of the increased temperatures”. John responds: “First, you can’t rule out the possibility that those CO<sub>2</sub> figures are just wrong. I, for one, don’t trust the IPCC’s figures. Secondly, I could accept everything you have just said, but it still wouldn’t *prove* that the man-made rise in CO<sub>2</sub> levels is causing the earth to warm up. You just can’t know that. It’s speculation.” Annoyed, Sally nevertheless doesn’t dispute what John says and the conversation moves on.

I suggest that in this case, John is attempting to raise the epistemic standard, and furthermore there is something inappropriate in his attempt to do so. The case bears similarities to other cases where the epistemic standard is raised, such as the bank case: notably, John raises a possible alternative that Sally cannot straightforwardly rule out, and then on that basis argues that Sally cannot know that the rise in CO<sub>2</sub> levels is causing the earth to warm up. I would also suggest that use of the word “prove” is often indicative of this kind of behaviour; it is an attempt to say to a person that their evidence is not sufficient given the standard that is in play at that point. Furthermore, Sally, although annoyed, does allow John’s more stringent use of ‘know’ by choosing not to challenge him. A great deal of work has been done on proposing theories of how the epistemic standard rises and falls, and I will not provide a complete account here. But I am committed to the claim that there are cases like *Climate change* where the epistemic standard is successfully raised, even if there is something indecent about it.

If we are going to diagnose John’s behaviour as in some sense indecent, then we face the challenge of identifying what makes it indecent. In particular, we have to provide some way of distinguishing *Climate change* from cases where the standard is raised but not inappropriately. The remainder of this section will be focused on sketching a way that this might be done.

One initial approach would be to draw on some of the recent work that has been done on conceptual engineering in order to capture John’s behaviour as a case of inappropriate conceptual engineering. Hansen (2019) distinguishes between successful and unsuccessful metalinguistic proposals – attempts to get speakers to adopt a proposed use. One way to be successful with a metalinguistic proposal is to hold the appropriate kind of linguistic authority. For example, “a journal editor could require an author to use ‘they’ and ‘their’ as gender-neutral pronouns instead of ‘he’ and ‘his’” (2019: 3). But one can also make metalinguistic proposals without any special authority if one is “backed by reasons that justify the speaker’s belief that using a particular expression in a particular way is a good idea” (2019: 3). So in order to instigate a change

in meaning, one must be backed by good reasons. It is useful to consider whether we can apply this idea to *Climate change* and argue that John lacks the requisite reasons. However, merely saying that John raises the standard for the wrong reason just leads to the question of what the good and bad reasons for raising the epistemic standard are. To answer this, we have to be more specific in thinking not just about the possibility of meaning change and metalinguistic proposals more generally, but of the possibility of changing our knowledge discourse. In this way, our focus shifts from the linguistic to the epistemic.

In order to consider ways that the epistemic standard could be inappropriately manipulated, it is worth considering what the function of a variable epistemic standard is. A full account of the function of a variable epistemic standard will necessarily appeal to what the point is of having any epistemic standard in the first place i.e. what the point is of having a knowledge threshold, rather just having evidence, credences, beliefs, etc. As alluded to earlier, there are a range of possible answers to this question (the role of knowledge in assertion, action etc.), and it is a substantial theoretical task to discern whether such answers can be viewed as consistent with one another. The answer I will focus on appeals simply to the value of knowledge. With knowledge comes value of a particular kind, and it is not clear that we could still possess that value in the same way if there was no distinction between knowledge and mere belief. Of course, this is only informative once we provide some account of what it is about knowledge that makes it valuable, but here I will rely upon some fairly general remarks without providing a detailed theory. We can distinguish between two kinds of value that knowledge arguably has. First, knowledge clearly has a practical value in being such that it can play a role in our practical deliberations, and so help us navigate our way around our environment in order to meet our aims and desires. If I want to get to Larissa, then my knowing how to get to Larissa will be of clear practical value. Second (and more controversially) knowledge arguably possesses a value that goes beyond its practical use, a kind of purely epistemic value. That is, we value knowledge purely in virtue of being the kind of inquisitive creatures that we are. Goldman sums up these two kinds of value well:

“All men by nature desire to know.” So said Aristotle (1924) and he was right. Information seeking is a pervasive activity of human life. We scan the horizon to see if rain is imminent, we watch the news to learn who has been elected, and we listen to the traffic report to anticipate delays in our drive to work. Our interest in information has two sources: curiosity and practical concerns. The dinosaur extinction fascinates us, although knowing its cause would have no material impact on our lives. We also seek knowledge for practical reasons, as when we solicit a physician’s diagnosis or compare prices at automobile dealerships. (Goldman 1999: 3)

So the value of knowledge – both practical and purely epistemic – gives us some explanation of why we should have an epistemic standard at all. But why a *variable* epistemic standard? It is notable that in work on epistemic standards (particularly in the contextualism and pragmatic encroachment literature), there are two recognised ways that the standard might rise. First, the standard might rise due to the high practical stakes involved in being right about the believed proposition. This is one way to interpret the bank case – the fact that it is much more important that Keith is right on whether the bank is open on Saturday raises the epistemic standard to a point where it can no longer be said that Keith does know the bank will be open. The other way is less concerned with practical stakes and more concerned with pursuing a more stringent form of inquiry. For example, the contextualist account of the sceptic is as someone who does

not raise the standard because of the practical stakes involved, but because they seek to make certain possibilities salient. What is striking is that these two ways that the epistemic standard are raised match up fairly well to the two kinds of value associated with knowledge that were outlined in the previous paragraph. Practical stakes raising clearly corresponds to the practical value possessed by knowledge, while stringent inquiry raising clearly corresponds to the (more controversial) purely epistemic value possessed by knowledge. What this suggests is that the variable epistemic standard is a way of better pursuing these kinds of value. On the one hand, in some contexts, the practical benefit in being right will be such that the practical value of knowledge will be enhanced by adjusting what it means to know. On the other hand, in some contexts, the purely epistemic benefit in pursuing a form of inquiry will be such that the epistemic value of knowledge will be enhanced by adjusting what it means to know. I suggest then, that there are only two very broad reasons to raise the epistemic standard. It is only permissible to raise the epistemic standard if:

- A more stringent epistemic standard would allow for a more practically valuable conception of knowledge, given the practical stakes involved, or
- A more stringent epistemic standard would allow for a more epistemically valuable conception of knowledge, given the nature of your inquiry.

The key claim of this paper, then, is that John in *Climate change* fails to meet this condition. He is attempting to raise the epistemic standard, not because he takes the practical stakes to call for it, nor does he genuinely want to pursue a more stringent form of inquiry, but because he is hoping to obfuscate, dominate an argument, and possibly even shut down inquiry. It is in this sense that we can label John's standard raising as inappropriate and we can distinguish from other cases – such as the bank case – where the standard is raised but not for inappropriate reason.

Earlier, I used Brister's work on the sceptic as bully as a motivation to look for more ordinary cases of inappropriate standard-raising, of which I count *Climate Change* as one instance. But isn't there an important difference here between the sceptic and John – isn't the sceptic attempting to engage in a more stringent form of inquiry, in precisely the way I have argued can lead to an acceptable form of standard-raising? If what I have said is right, there is nothing wrong with the sceptic's raising of the standard per se provided that they do so in the pursuit of a more stringent form of epistemic inquiry. I think that is right and would help capture the fact that sometimes – in academia or elsewhere – we do pursue a form of inquiry with a higher standard, if only for our own curiosity. Provided that this is done to pursue the epistemic value of knowledge, I argue that there is nothing wrong in doing so. The same point applies to John. The explanandum of the paper has been the apparent wrongness of John's behaviour. I have argued that his behaviour seems wrong insofar as it is natural to interpret his behaviour as motivated by the wrong kind of reasons. But John's reasons for acting as he does were not explicitly included as part of the case, so as it stands it is consistent with my account that John is in fact raising the standard in an acceptable way.<sup>12</sup> Strictly then, my account should only be committed to the following conditional: if John is raising the epistemic standard for reasons other than those outlined by my account, then he is acting in an epistemically wrongful way. It is this conditional, and the fact that we naturally interpret John as making the antecedent true, that explains the apparent wrongfulness of the case. But it is consistent with my account that there may be nothing

<sup>12</sup>Thanks to an anonymous referee for emphasising this point.

wrong with being a climate change denier per se just as there may be nothing wrong with being an external world sceptic.

My account is attractive insofar as it captures similar cases that have been previously discussed in the literature. Blome-Tillman (2014: 19) discusses a humorous case where a teenage boy is caught sneaking back home in the middle of the night by his parents. Imagine that the boy responds to being caught by saying “You can’t know that I was sneaking in just now. After all, you may have dreamt that, or you may be a brain in a vat”. Blome-Tillman’s point is that this is an unsuccessful attempt at raising the epistemic standard provided that the parents fail to take the possibilities that the boy raises seriously. But suppose that the boy is successful in raising the epistemic standard and the parents subsequently ruminate on the possibility of being a brain in a vat rather than admonishing their son. There still seems to be something epistemically inappropriate in what the boy has done. This is precisely what my account predicts: the boy raised the epistemic standard for unacceptable reasons – he did it to avoid punishment rather than to pursue the value of knowledge.

Thus far I have only tried to find a way to negatively evaluate John’s standard-raising in *Climate change*. I haven’t gone so far as to discuss the possibility that there is any kind of epistemic harm or injustice involved in the case, and indeed, this is not something I will do in any great detail in this paper. That said, it is worth noticing that there is a strong similarity between John’s behaviour and the notion of *hermeneutical injustice* as introduced by Fricker (2007). Hermeneutical injustice occurs when there is a lacuna in the collective epistemic resources that prevents the proper understanding of a given phenomenon. Most clearly, this occurs when a society lacks some concept that would be required in order to understand a particular kind of experience. One of Fricker’s central cases is that in which a society lacked the concept of sexual assault, and so women and others who suffered experiences that fall under that concept were unable to properly understand or communicate about their experience. There is a sense in which John is attempting to instil a kind of localised hermeneutical injustice. Inappropriate standard raisers such as John follow the tactic of the sceptic by raising the epistemic standard to such a level that the distinction between knowers and non-knowers no longer cuts interestingly across our epistemic community. In doing so, the standard raiser robs us of the distinction between knowers and non-knowers that would otherwise have been in play had the epistemic standard not been raised. Now clearly, this isn’t hermeneutical injustice in the usual sense insofar as we regain access to that distinction once the epistemic standard returns to a normal level. But in that instant, in that context, we are momentarily prevented from using knowledge terminology to properly describe one another’s epistemic positions regarding the proposition at hand. In that sense, we can view the behaviour of John and other standard-raisers as a kind of fleeting, micro-form of hermeneutical injustice.<sup>13</sup>

I will finish by considering an objection. It could be argued that it is not possible for an instance of standard-raising to be inappropriate in the sense that I have outlined simply because the epistemic standard is always where it ought to be. Robin McKenna (2013) argues for a contextualist view of this kind, according to which “S knows p” is true in c only if S can rule out the alternatives that *ought to be* considered given the practical stakes involved in c. According to this normatively loaded view, it

<sup>13</sup>Fricker (2007: 22 ff.) is less inclined to describe a harm as an injustice if it is of a lesser severity and she is also primarily concerned with systematic rather than incidental cases. It may be for these reasons that the phenomenon of wrongful standard-raising does not constitute a form of hermeneutical injustice. My point is simply here to highlight the important similarity between the two phenomena – that both are cases where an agent lacks a conceptual resource and that this is to their detriment.

would either have to be the case that John successfully raises the epistemic standard and does so appropriately, or that he attempts and fails to raise the epistemic standard. I am not entirely averse to this alternative account of John as attempting and failing to raise the epistemic standards, and I think that much of what I have said in this paper could be accommodated in such a view. However, I prefer a view of the epistemic standard as one that is primarily sensitive to conversational manoeuvres within a discourse of inquiry, such as the making salient of certain alternative possibilities and certain practical stakes. In that sense, I prefer to adopt a view of the epistemic standard such that it can be manipulated by nefarious actors. This better fits with the wider picture from which conceptual engineering has been encouraged, according to which it is an open question whether we are using the concepts that are best suited to our purposes.

#### 4. Conclusion

In this paper, I have identified wrongful raisings of the epistemic standard as nefarious attempts at contextually restricted conceptual engineering regarding our knowledge discourse. Furthermore, I have provided a partial theory of what it is about such cases that make them inappropriate, and in doing so I have argued that there are two broad reasons to raise the epistemic standard and that inappropriate standard-raisings (such as *Climate change*) are characterized by the fact that neither reason is acted upon.

To end, it is worth returning to where we started. In discussing cases where the standard is raised wrongfully, I have sought to employ a more neutral terminology so that the conclusions of this paper are consistent with a range of background views on the epistemic standard (e.g. contextualism, pragmatic encroachment views, etc.). The possibility of using less theoretically committed claims such as modest variability has been largely overlooked in discussion of the epistemic standard, and this leads to the worry that epistemological progress on the epistemic standard and the role it plays is stifled until difficult issues surrounding semantics, pragmatics, and theory of knowledge are settled. In this paper, I hope to have shown how further insights can be shed on the epistemic standard in such a way that is not beholden to such controversial issues.<sup>14</sup>

#### References

- Aristotle.** (1924). *Metaphysics*, trans. W.D. Ross. <http://classics.mit.edu/Aristotle/metaphysics.1.i.html>.
- Blome-Tillmann M.** (2014). *Knowledge and Presuppositions*. Oxford: Oxford University Press.
- Briester E.** (2009). 'Feminist Epistemology, Contextualism, and Philosophical Skepticism.' *Metaphilosophy* 40(5), 671–88.
- Briester E.** (2017). 'Feminism and Contextualism.' In J.J. Ichikawa (ed.), *The Routledge Handbook of Epistemic Contextualism*, pp. 57–68. Abingdon: Routledge.
- Burgess A. and Plunkett D.** (2013a). 'Conceptual Ethics I.' *Philosophy Compass* 8(12), 1091–101.
- Burgess A. and Plunkett D.** (2013b). 'Conceptual Ethics II.' *Philosophy Compass* 8(12), 1102–110.
- Cappelen H.** (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Carnap R.** (1950). *Logical Foundations of Probability*. Chicago, IL: Chicago University Press.
- Chomsky N.** (1977). *Essays on Form and Interpretation*. New York, NY: Elsevier North-Holland.
- Cohen S.** (1999). 'Contextualism, Skepticism, and the Structure of Reasons.' *Noûs* 33(s13), 57–89.
- DeRose K.** (1992). 'Contextualism and Knowledge Attributions.' *Philosophy and Phenomenological Research* 52(4), 913–29.
- Fantl J. and McGrath M.** (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

<sup>14</sup>This paper has been greatly improved by comments from Corine Besson, the audience at the Post-Truth and the Objectivity of Epistemic Norms Conference at the University of Sussex, and two anonymous referees.



- Fricker M.** (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Gao J., Gerken M. and Ryan S.B.** (2017). 'Does Contextualism Hinge on A Methodological Dispute?' In J.J. Ichikawa (ed.), *The Routledge Handbook of Epistemic Contextualism*, pp. 81–93. Abingdon: Routledge.
- Gerken M.** (2013). 'Epistemic Focal Bias.' *Australasian Journal of Philosophy* **91**(1), 41–61.
- Gerken M.** (2019). 'Pragmatic Encroachment and the Challenge from Epistemic Injustice.' *Philosophers' Imprint* **19**.
- Goldman A.** (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.
- Grindrod J.** (2019). 'Depth, Value, and Context.' *Ergo* **6**(24). <https://doi.org/10.3998/ergo.12405314.0006.024>.
- Grindrod J., Andow J. and Hansen N.** (2018). 'Third-person Knowledge Ascriptions: A Crucial Experiment for Contextualism.' *Mind and Language*. doi: 10.1111/mila.12196.
- Hansen N.** (2019). 'Metalinguistic Proposals.' *Inquiry*. <https://doi.org/10.1080/0020174X.2019.1658628>.
- Hansen N. and Chemla E.** (2013). 'Experimenting on Contextualism.' *Mind and Language* **28**(3), 286–321.
- Haslanger S.** (1999). 'What Knowledge is and What it Ought to be: Feminist Values and Normative Epistemology.' *Philosophical Perspectives* **13**(Epistemology), 459–80.
- Haslanger S.** (2000). 'Gender and Race: (What) Are They? (What) Do We Want Them To Be?' *Noûs* **34**(1), 31–55.
- Hawthorne J.** (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Ludlow P.** (2014). *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.
- MacFarlane J.** (2011). 'Relativism and Knowledge Attributions.' In D. Pritchard and S. Bernecker (eds), *The Routledge Companion to Epistemology*, pp. 536–44. London: Routledge.
- McKenna R.** (2013). 'Epistemic Contextualism: A Normative Approach.' *Pacific Philosophical Quarterly* **94**(1), 101–23.
- Nagel J.** (2010). 'Knowledge Ascriptions and the Psychological Consequences of Thinking About Error.' *Philosophical Quarterly* **60**(239), 286–306.
- Plunkett D. and Sundell T.** (2013). 'Disagreement and the Semantics of Normative and Evaluative Terms.' *Philosophers' Imprint* **13**(23), 1–37.
- Pritchard D.** (2012). 'On Meta-Epistemology.' *Harvard Review of Philosophy* **18**(1), 91–108.
- Rysiew P.** (2001). 'The Context-Sensitivity of Knowledge Attributions.' *Noûs* **35**(4), 477–514.
- Rysiew P.** (2005). 'Contesting Contextualism.' *Grazer Philosophische Studien* **69**(1), 51–70.
- Stanley J.** (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Williamson T.** (2005). 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge.' *Philosophical Quarterly* **55**(219), 213–35.

**Jumbly Grindrod** is a Lecturer at the University of Reading. His research interests lie in epistemology and philosophy of language.