

This is the penultimate draft of a paper for *Skeptical Invariantism Reconsidered*, which is being edited by C. Kyriacou and K. Wallbridge, and coming out with Routledge

## **A (Partial) Defence of Moderate Skeptical Invariantism**

Robin McKenna

University of Liverpool

[rnmckenna@gmail.com](mailto:rnmckenna@gmail.com)

**Abstract** Skeptical invariantism isn't a popular view about the semantics of knowledge attributions. But what, exactly, is wrong with it? The basic problem is that it seems to run foul of the fact that we know quite a lot of things. I agree that it is a key desideratum for an account of knowledge that it accommodate the fact that we know a lot of things. But what sorts of things should a plausible theory of knowledge say that we know? In this paper I sketch an answer to this question and then apply it to skeptical invariantism. I start by distinguishing between a "radical" skeptical invariantist position (on which the standards for knowing are so high that we rarely if ever satisfy them) and a more "moderate" skeptical invariantist position (on which the standards are very hard, but not impossible, to satisfy). I then argue that, while radical forms of skeptical invariantism are clearly not going to do a good job of satisfying the key desideratum, more moderate forms can do quite well with respect to it. In particular, I argue the version of moderate skeptical invariantism defended by Davide Fassio can plausibly satisfy it.

### **0. Introductory Remarks**

Skepticism is not a popular view in epistemology. Most epistemologists reject it.<sup>1</sup> Why? Some reject it because they view the arguments for it as based on a mistake (Williams 1991; Wright 1991; Sosa 2009). Others reject it for the simple reason that they think we know lots of things (Moore 1939; Pritchard 2002; Sosa 1999). These epistemologists take it as a key desideratum on an account of knowledge that it fits with the (pre-theoretic?) intuition that we know a lot of things. Because skepticism, by definition, cannot accommodate this desideratum, it is to be rejected.

I think that something like this desideratum is right. We do know certain things, and any view of knowledge which says we don't is therefore to be rejected. This might not sound like a promising starting point for an exploration of skeptical invariantism. This paper looks like it is over before it has started.

But this would be too hasty for three reasons. The first is that the desideratum is vague. *Which* things do we know? It makes no sense to require an account of knowledge to vindicate every intuition someone might have about the extension of

---

<sup>1</sup> Exceptions include Fumerton (1995), Stroud (1984) and Unger (1975).

the concept of knowledge. We need to come up with, and defend, a list of things we should want an account of knowledge to say we know.

Second, I am open to *revisionary* accounts of knowledge (Craig 1990; Fassio and McKenna 2015; Haslanger 1999; Kusch and McKenna forthcoming). An account of knowledge that is theoretically attractive may also revise certain aspects of our ordinary understanding of knowledge, such as how much of it we actually have. But there has to be a trade-off here: the more revisionary an account is, the greater the theoretical benefits must be (Fassio and McKenna 2015).<sup>2</sup>

Third, the label “skeptical invariantism” denotes a family of views (Cappelen 2005; Davis 2007; Fassio forthcoming; Kyriacou 2017; Levin 2008). Some of these views are “radical” in that they hold that we know very little, if anything. But others are more “moderate” in that they merely hold that we know significantly less than we ordinarily take ourselves to know. You might hold that we know a lot less than we think we do, yet still hold that we know a fair bit.

Putting all this together, it should be clear that, properly understood, the key desideratum need not rule out skeptical invariantist accounts of knowledge from the outset. But is there a version of skeptical invariantism that can actually accommodate it? This is my question in this paper.

Here is the plan. In the first part I distinguish between “radical” and “moderate” skeptical invariantism, and identify specific versions of each view that I focus on in the rest of the paper.

In the second part I say more about what I take the core cases of knowledge to be. Over and above the sorts of cases you might expect—e.g. simple perceptual beliefs formed in normal conditions, knowledge of a priori, analytic or necessary truths—I argue we should follow the naturalised epistemologist and look to science to identify core cases of knowledge. I focus on one item of scientific knowledge in particular: we know that global temperatures are rising.

In the third part I investigate whether skeptical invariantism can accommodate the core cases identified in the second part. I focus on Davide Fassio’s (forthcoming) moderate skeptical invariantism and investigate how it fares with the key

---

<sup>2</sup> Michael Hannon suggested that it is unclear whether radical skepticism really is revisionary, because the radical skeptic generally does *not* think we should change our ordinary practice of knowledge ascriptions (see e.g. Unger (1975)). This is a fair point, but it does seem to require a workable pragmatic account of knowledge ascriptions. For discussion of skeptical invariantism and the pragmatics of knowledge ascriptions see Cappelen (2005), Davis (2007), (DeRose 2017), Dinges (2016) and Kyriacou (2017).

desideratum. While this investigation raises several issues, I argue that his view can go some way towards accommodating the cases, and so satisfying the desideratum. This by no means constitutes a full defence of moderate skeptical invariantism. But it does show that it is a view well worth taking seriously. Maybe we really do know a lot less than we ordinarily take ourselves to know.

## 1. Skeptical Invariantism

Since the late 1980s there has been much debate in epistemology about the semantics of knowledge ascriptions—utterances of sentences including the word “knows” and its cognates. The crucial question has been: is the word “knows” (and its cognates) context-sensitive, in the same way that indexicals (“I”, “here”) or gradable adjectives (“tall”, “heavy”) are context-sensitive? That is, do uses of the word “knows” refer to different relations in different contexts of utterance? The “invariantist” holds that they refer to the same relation in all contexts of utterance: the knowledge relation (Brown 2006; Hazlett 2009; Rysiew 2001). In contrast, the “contextualist” holds that they refer to different relations in different contexts of utterance (DeRose 1992; Cohen 1986; Lewis 1996).

While more could be said about the differences between invariantism and contextualism, we can put contextualism to one side and focus on invariantism. There are two important distinctions to draw between different forms of invariantism. The first is between “traditional” and “interest-relative” invariantism. The traditional invariantist and the interest-relative invariantist agree that the word “knows” and its cognates is context-insensitive: uses of it express the same relation irrespective of context. They differ over the nature of that relation. For the traditional invariantist, whether someone knows something depends on purely epistemic or truth-connected factors, such as the strength of their evidence (Brown 2008; Gerken 2017; Reed 2010). For the interest-relative invariantist, whether someone knows depends on a combination of epistemic and “practical” factors, such as the costs of being wrong or “stakes” (Fantl and McGrath 2009; Hawthorne 2004; Stanley 2005).

Imagine Catriona and Laurie both have the same evidence that the next stop is South Parkway, and they both believe that the next stop is South Parkway on the basis of this evidence. For Catriona it doesn’t matter if she is wrong: she doesn’t need to get off at South Parkway, as she can equally well get off at the next stop. In contrast, for Laurie it is absolutely imperative that she not get this wrong. For the traditional invariantist, this difference in what is “at stake” is irrelevant: either the evidence is sufficient for them both to know, or it isn’t. For the interest-relative invariantist, the difference in stakes matters. It may be that the evidence is sufficient for Catriona to know given how little it would matter if she were wrong, but it is insufficient for Laurie to know, given how much it would matter if she were wrong.

The second distinction is between “skeptical” and “moderate” invariantism. The invariantist holds that the word “knows” is univocal: it expresses the same relation in all contexts of utterance. But how (to put it roughly) *stringent* is this relation? That is, how high are the standards for knowledge? The moderate invariantist says that the standards are relatively modest.<sup>3</sup> They are high enough to draw a distinction between, say, merely truly believing that it is raining and knowing that it is raining. But they are sufficiently low that we count as knowing an awful lot. The skeptical invariantist says that the standards are considerably higher—high enough that we count as knowing a lot less than the moderate invariantist thinks we do (Cappelen 2005; Davis 2007; Fassio forthcoming; Kyriacou 2017; Levin 2008).

With these distinctions in hand, we can now focus on skeptical invariantism. I start with two preliminaries. First, as I use the label, skeptical invariantism is neutral between traditional and interest-relative invariantism.<sup>4</sup> That is, the skeptical invariantist might hold that knowledge depends purely on epistemic factors, but she might also allow that knowledge depends partly on pragmatic factors.

Second, skeptical invariantism denotes a family of positions, ranging from a view on which the standards for knowledge are very high indeed, so high that we rarely if ever meet them, to views on which the standards for knowledge are more stringent than the moderate invariantist would allow, but not so high that we can’t sometimes meet them. We can distinguish between “radical” and “moderate” skeptical invariantism. But these are perhaps best seen as labels for two contrasting tendencies, rather than for distinct views: the radical skeptical invariantist tends to insist that the standards for knowledge are about as high as they could be, whereas the moderate skeptical invariantist tends to insist that the standards are very high, but not so high as the radical skeptical invariantist insists they are.

In the rest of the paper it will be helpful to have concrete examples of radical and moderate skeptical invariantist positions to hand. I focus on the radical skeptical invariantism defended by Christos Kyriacou (2017) and the moderate skeptical invariantism defended by Davide Fassio (forthcoming), though for reasons that

---

<sup>3</sup> The label “moderate invariantism” is from Hawthorne (2004). All of the invariantists cited above are moderate in Hawthorne’s sense.

<sup>4</sup> I make this stipulation because I don’t want to get embroiled in discussion about whether the moderate skeptical invariantist view I go on to consider (Fassio forthcoming) is really a version of traditional invariantism. It agrees with the traditional invariantist that, if a subject knows in a low stakes case, then she knows in the corresponding high stakes case. But it characterises knowledge in terms of a notion (practical certainty) that looks a little pragmatic. Part of the issue here is that the terms “epistemic” and “pragmatic” are not usually clearly defined and distinguished. For discussion see Hannon (2020).

will become clear below I focus more on Fassio's moderate skeptical invariantism than Kyriacou's radical skeptical invariantism.

Kyriacou defends what he calls "bifurcated skeptical invariantism". His skeptical invariantism is "bifurcated" because he holds that there are two related but distinct concepts of knowledge, both of which are operative, but in different contexts. These are (put roughly) a concept of knowledge as fallible and a concept of knowledge as infallible. These concepts are related in that they hold that knowledge requires justification. But they are distinct in that they take different views of the nature of justification. In order to have fallible knowledge, one must merely believe on the basis of strong but defeasible justification. In contrast, in order to have infallible knowledge, one must believe on the basis of a justification that entails the truth of one's belief.<sup>5</sup>

Kyriacou's view is that, so far as the semantics of "knowledge" and knowledge ascriptions is concerned, to say that someone "knows" something is to say that they have infallible knowledge. Thus, many of our knowledge ascriptions—indeed, all of the ones where the subject lacks infallible justification—are, strictly speaking, false.<sup>6</sup> But, so far as the pragmatics of "knowledge" and knowledge ascriptions is concerned, we often use them to communicate that subjects have fallible knowledge. Because subjects often do have fallible knowledge—they often believe on the basis of fallible justification—our knowledge ascriptions often pragmatically convey truths, even though they are largely, strictly speaking, false.

Now turning to Fassio, there are two crucial components of his view:

[A] subject knows  $p$  only if she is *practically certain* that  $p$ , where practical certainty can be defined as the confidence a rational subject would have to have for her to believe that  $p$  and act on  $p$  no matter the stakes on being right about whether  $p$  (forthcoming, 3).

S is practically certain that  $p$  if and only if [if S were rationally responsive to stakes, S would maintain a degree of confidence sufficient to believe that  $p$  and act on  $p$  no matter how much turned on  $p$ ] (forthcoming, 5).

---

<sup>5</sup> In contemporary epistemology infallibilism is enjoying a renaissance, in large part due to Williamson (2000). Williamson identifies one's evidence with one's knowledge so, when you know that  $p$ , you have evidence which entails  $p$  (i.e. your knowledge that  $p$ ). Kyriacou thinks his view entails that we know little, if anything, so he clearly must reject Williamsonian infallibilism. More generally, he must reject any view on which our evidence goes beyond our experiences and includes claims about the external world. For discussion of Williamsonian infallibilism see Brown (2018).

<sup>6</sup> Recall from fn. 5 that I am ignoring Williamsonian infallibilism.

The thought is that you know that  $p$  only if you are practically certain that  $p$ , and you are practically certain that  $p$  iff you are sufficiently confident that  $p$  such that, no matter the stakes, the rational response to them would be to still believe that  $p$  and act on  $p$ . Note that Fassio does not “build it in” to his view that we know more than the radical skeptical invariantist would allow, but less than the moderate invariantist would allow.<sup>7</sup> Rather, he provides us with a condition on knowledge (practical certainty), which he argues is easier to meet than the sort of condition a radical skeptical invariantist would endorse, but harder to meet than the sort of condition a moderate invariantist would endorse.

To see how his “practical certainty condition” works, consider one of the cases above. Recall that Catriona and Laurie had the same evidence that the next stop is South Parkway, and they both believed that the next stop is South Parkway on the basis of this evidence. To determine whether either of them have knowledge we need to ask if they are practically certain that the next stop is South Parkway, and they are practically certain that the next stop is South Parkway if and only if, if they rationally responded to the stakes, they would be confident enough to believe and act on the proposition that the next stop is South Parkway no matter the stakes.<sup>8</sup>

Whether this condition is met is going to depend on the strength of the evidence. We can distinguish three cases. First, let’s imagine it is relatively flimsy: they both glanced at the list of stops on the platform before they boarded the train, and they seem to recall that South Parkway was on it. In this case, it is clear that the condition will not be met. Imagine a high stakes situation. In such a situation, if they rationally responded to the stakes, they would be insufficiently confident to believe that the train stops at South Parkway, or to act on this basis. So they aren’t practically certain.

Second, let’s imagine it is considerably stronger, but still far from conclusive: they have taken this train before, and it usually stops at South Parkway. In this case, it is (or at least Fassio thinks it is) clear that the condition will also not be met. That is,

---

<sup>7</sup> Compare: a safety theorist like Sosa (2007) does not built it into his account that we know anything, but rather proposes a particular condition on knowledge (safety), then argues that many of our beliefs are safe in his sense.

<sup>8</sup> On Fassio’s view, what matters is what the rational response to the stakes would be, not what some subject’s actual response would be. You might ask: what is it to rationally respond to the stakes? For Fassio’s discussion of this see Section 1 of his paper. But the rationale for the rationality requirement is that Fassio wants to avoid situations where a subject would remain confident enough to believe something, or act on the basis of it, no matter the stakes because they are habitually over-confident. His thought is that the over-confident subject would not be rationally responding to the stakes,

we can imagine a situation where the stakes are high enough that, if they were to rationally respond to them, they would be insufficiently confident to believe that the train stops at South Parkway, or to act on this basis. So they aren't practically certain here either.

Finally, let's imagine it is very strong: they just heard an announcement from the driver saying the next stop is South Parkway, all the signs say the next stop is South Parkway, they recognise that they are approaching South Parkway, the train is slowing down, and so on. Fassio's view is that, in this case, the condition *is* met. That is, he thinks that, no matter how high we imagine the stakes being, it remains the case that, if they were to rationally respond to them, they would be confident enough to believe that the train stops at South Parkway, and to act on this basis. So they are practically certain.

These three cases not only illustrate how the practical certainty condition works; they also illustrate the "moderate" and "skeptical" aspects of Fassio's view. The view is moderate in that, in the third case, there plausibly is knowledge. The view is skeptical in that, in the second case, there isn't. Thus, they distinguish Fassio's view from moderate invariantism (which would hold there is knowledge in the second case), and radical skeptical invariantism (which would deny there is knowledge in the third case).

This completes my overview of skeptical invariantism. I am now going to consider the extent of our knowledge, and whether there is a skeptical invariantist position that can plausibly accommodate the things we should take ourselves to know.

## **2. Articulating the Key Desideratum**

How much do we know? And, more importantly for our purposes, how much does a theory of knowledge need to say we know in order to be plausible? In this section I want to make two suggestions. The first—which I won't spend much time on—is that the debate about invariantism and the semantics of knowledge ascriptions tends to implicitly assume an implausible answer to the second of these questions. The second—which I will spend a little more time on—is that we can draw up a list of several classes of beliefs that we should expect a theory of knowledge to vindicate as cases of knowledge. I should say from the outset that I make no claim to the effect that my list is complete. It may be there are classes of beliefs that a theory of knowledge should vindicate that are not on my list, and I won't discuss the extent to which the skeptical invariantist can accommodate them.

But my aim in this paper is to develop a partial, not a full, defence of (moderate) skeptical invariantism.<sup>9</sup>

Let's start with the first suggestion. In the debate about knowledge ascriptions it often seems to be taken as obvious that the subjects in the "low stakes" version of DeRose's bank cases have knowledge. This case goes like this:

*Low:* 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 a check. It's not important that they do so, as they have no impending bills. But, as they drive past the bank, they notice that the lines inside are very long. Realising that it isn't very important that the check is deposited right away, Hannah says, "I know that the bank will be open tomorrow. I was there just two weeks ago on Saturday morning. So we can deposit our check tomorrow morning".<sup>10</sup>

The general assumption seems to be that, clearly, Hannah speaks truly when she says that she knows the bank will be open on Saturday. Some views seek to vindicate this assumption by complicating it: for the contextualist, the subjects in these cases can truly be said to "know" in some contexts but not others, and for the interest-relative invariantist, they know in part because the stakes are low. Other views seek to vindicate it in a more straightforward manner (e.g. traditional moderate invariantism). The problem for the skeptical invariantist (whether radical or moderate) is meant to be that they deny that the subjects in these cases have knowledge irrespective of the stakes, and no matter in what context one ascribes knowledge to them. But, for this to be a problem, we need to be justified in treating low stakes bank cases as core or non-negotiable cases of knowledge, and this strikes me as highly implausible. Note: my claim is *not* that Hannah does not know in this case. Maybe she does. It may well be that our best account of knowledge rules that the subjects in bank cases have knowledge. That doesn't mean that we should build it into our theorising about knowledge that an acceptable account of knowledge has to have this result (cf. Dinges 2016)

Turning to the second suggestion, we can start with what Fassio says we can know according to his moderate skeptical invariantism. Fassio says:

---

<sup>9</sup> In personal correspondence Fassio told me that he thinks his account can also accommodate the following classes of knowledge: inferential knowledge (where inference need not be deductive), knowledge about what is probable or will be the case conditional on other events, modal knowledge, aesthetic knowledge, historical knowledge, medical knowledge and technical knowledge. If he is right about this, then perhaps a full defence of moderate skeptical invariantism can be mounted. But I lack the space to deal with all these classes of knowledge here.

<sup>10</sup> This version of the cases is from Stanley (2005, 3). For the original case see DeRose (1992).

MSI [moderate skeptical invariantism] allows that knowledge is compatible with absence of absolute certainty. One may be practically certain, and therefore in a position to know that  $p$ , even if one is not in a position to exclude some abstract error possibilities. While MSI excludes knowledge in Bank-like cases (a subject that is rational and aware of the high stakes will lack the confidence necessary to believe and act on the proposition that the bank is open on Saturday), all beliefs based on a rational degree of confidence that is ‘robust’ enough to survive any significant increase in the stakes are good candidates for knowledge (forthcoming, 9).

But what beliefs are based on a rational degree of confidence robust enough to survive any increase in the stakes? Fassio is reluctant to list concrete examples, but he ventures the following:

- His belief that the capital of Italy is Rome.
- His belief (formed on a Monday) that tomorrow is Tuesday.
- His belief that he ate spaghetti for lunch.
- His belief that in December he bought presents for his family.
- His belief that there are two oranges on the table.
- Certain testimonial beliefs (e.g. Hannah’s belief that the next stop is South Parkway, based on being told by a ticket inspector a few minutes before).

Rather than focus on concrete examples, I think a better approach is to focus on broad *classes* of beliefs which, on a plausible account of knowledge, we should be said to know. The claim is not that any particular individual has some piece of knowledge. Rather, the claim is that many subjects in certain sorts of situations know certain things. For instance:

1. Beliefs about our immediate surroundings formed on the basis of perception in normal conditions (“there are two oranges on the table”).
2. Very well supported memorial and testimonial beliefs (“I ate spaghetti for lunch”, “the next stop is South Parkway”).
3. Beliefs in matters of fact that we generally take for granted, including facts about the future (“the capital of Italy is Rome”, “tomorrow is Tuesday”, “the sun will rise tomorrow”).

In addition to these three classes of beliefs suggested by Fassio, I would like to add some more. Consider:

4. Some introspective beliefs (e.g. “I am in pain right now”).

5. Beliefs in a priori/necessary/analytic truths (e.g. “all bachelors are unmarried”, “ $2+2=4$ ”, “water is  $H_2O$ ”).<sup>11</sup>

Finally, I want to add certain scientific beliefs. Now, I don’t think that an account of knowledge should show that we know all and everything that science takes us to know. Scientists may differ as to what we know, and it may be that some scientists take us to know that some hypothesis  $H$  is true whereas others take us to know that  $H$  is false. Rather, I think that, when it comes to hypotheses that are established with as much certainty as science can provide us with, an account of knowledge should allow that we do indeed know that these hypotheses are true. I want to focus on a particular example here: science, I take it, has shown with as much certainty as we could hope for that global temperatures are rising due to human activities. Of course, science hasn’t established that global warming is happening due to human activities with absolute certainty. But science doesn’t deal in absolute certainties.<sup>12</sup>

I offer two reasons for including certain items of scientific knowledge in our key desideratum, the first of which is less theoretically loaded than the second.

First, it would be more than a little odd to accept testimonial beliefs, like Catriona’s belief that the train stops at South Parkway, as core cases of knowledge, yet not to accept our belief that global warming is happening due to human activities as a core case. I don’t know about you, but insofar as the comparison makes sense, we have far more and far better evidence that human activities are causing global temperatures to rise than Catriona does that the next stop is South Parkway.<sup>13</sup> Similar points can, I think, be made about memorial beliefs, and perhaps some perceptual beliefs.

---

<sup>11</sup> On some views we can distinguish between a priori, necessary or analytic truths and on some views we can’t. For my purposes it doesn’t much matter where you stand on these issues.

<sup>12</sup> While it is hard to quantify the degree of certainty, we do know the degree of consensus among climate scientists that human activities are causing global warming, which is 97% (see Cook et al. 2016). Degree of consensus is, of course, a poor proxy for degree of certainty, but it is worth noting that such a high degree of consensus is unusual.

<sup>13</sup> Of course, I’m comparing a case of individual testimonial knowledge with group scientific knowledge, so there are some differences. If you like, replace talk of “us” knowing that global warming is happening due to human activities with talk of individual climate scientists knowing this. I still think it is right to say that an individual climate scientist has better evidence that climate change is happening due to human activities than Catriona has from testimony in the case we’re imagining.

One might object that what science has established is not that climate change is happening due to human activities, but rather that it is (very) likely that climate change is happening due to human activities, and so what we can be said to know is not that climate change is happening, but that it is (very) likely that it is.<sup>14</sup> Of course, science does establish the likelihood claim. But why not also take it to establish the unqualified claim? If you agree that we have more and better evidence for some scientific hypotheses than we do for testimonial (and memorial and some perceptual) beliefs, then it would be odd to take science to only establish claims about probabilities, and yet to take it that we can come to know what the next step is on the basis of testimony, rather than merely that it is (very) likely that the next step is such-and-such.

Second, a key insight from naturalised epistemology is that, when it comes to the traditional epistemological question of the extent of our knowledge, we should look to science to answer it (cf. Quine 1969). Now, you might object that there are things we clearly seem to know that have nothing to do with science. Take, for instance, simple perceptual knowledge or memorial knowledge. These forms of knowledge can be scientific, but they need not be. This point is fair so far as it goes but, first, you might respond that science can explain the mechanisms involved in producing perceptual and memorial knowledge and, second, this would only show that we can't look to science to identify the *full* extent of our knowledge.

I have identified six classes of things we want an account of knowledge to say we know:

1. Beliefs about our immediate surroundings formed on the basis of perception in normal conditions.
2. Very well supported memorial and testimonial beliefs.
3. Beliefs in matters of fact that we generally take for granted.
4. Some introspective beliefs.
5. A priori/necessary/analytic truths.
6. Some items of scientific knowledge (e.g. “global temperatures are rising due to human activities”).

I now turn to seeing how skeptical invariantism, particularly moderate forms of it, does with respect to this desideratum.

### **3. Assessing Skeptical Invariantism**

I have argued that it is a key desideratum on an account of knowledge that it accommodates certain core cases of knowledge. While for obvious reasons I will

---

<sup>14</sup> Thanks to Christos Kyriacou and Michael Hannon for pushing me to say something about this objection.

focus on moderate skeptical invariantism, I want to—albeit very briefly—say something about the basis on which I am dismissing radical forms of skeptical invariantism, and particularly Kyriacou’s bifurcated skeptical invariantism.

I myself have no problem in principle with revisionary accounts, whether of knowledge or of anything else. A revisionary account of knowledge may be more theoretically attractive than its rivals, and this may provide ample reason for revising our ordinary understanding of knowledge, or of its extent (see Fassio and McKenna (2015)). But I do think that there are certain constraints on any revisionary project. In Fassio and McKenna (2015) we argue for a “moderately conservative” revisionary methodology, according to which, the more revisionary an account of knowledge is, the higher the barrier to accepting it. Put simply, an account that radically revises our ordinary understanding of knowledge—for instance, by claiming that we know an awful lot less than we ordinarily take ourselves to know—would not just have to secure several theoretical benefits. It would have to secure far more theoretical benefits than any rival account.<sup>15</sup> While Kyriacou claims some theoretical benefits for his radical skeptical invariantist account of knowledge (e.g. it can (dis)solve the Gettier problem, and it addresses Kripke’s dogmatism paradox), I don’t think he makes (or attempts to make) the case that it secures far more theoretical benefits than any rival account.

While more could be said here, I now turn to my main question in this section: can Fassio’s moderate skeptical invariantism accommodate the core cases of knowledge identified in the previous section? While I won’t give an unqualified affirmative answer to this question, it is at least initially plausible that it can.

We can start with perceptual, memorial and testimonial beliefs. Fassio spends some time arguing that we can be practically certain when it comes to some of our perceptual, memorial and testimonial beliefs. We can split his argument into two parts. First, he considers an objection which, if right, would show that none of our perceptual, memorial or testimonial beliefs could ever be practically certain.<sup>16</sup> Recall Catriona. Can she be practically certain that the next stop is South Parkway, on the basis of being told that it is by the ticket inspector a few moments before? Consider a context where Catriona is offered an extremely high stakes bet with the following pay-off structure. If Catriona is right that the next stop is South Parkway, she gets a penny. If Catriona is wrong about this, something terrible will happen to her—make this thing as terrible as you can imagine. Many think that, in such a

---

<sup>15</sup> This is perhaps where the defence of the justified true belief account of knowledge considered (though not endorsed) in Weatherson (2003) goes wrong. There are *some* theoretical benefits to holding on to the justified true belief account, but they are surely not sufficient to adequately defend it.

<sup>16</sup> For this sort of objection see Weatherson (2012).

context, if Catriona were to rationally respond to the stakes, she would not be confident enough that the next stop is South Parkway to take the bet. Indeed, the rational thing to do would be to decline to take the bet.

Fassio's response to this strikes me as right. He agrees that, in this sort of case, it would not be rational for Catriona to take the bet. But he denies that this poses a problem for his account of knowledge. Now, "bet cases" would be a problem if the reason why it would not be rational for Catriona to take the bet is that the stakes are so high. But, thinks Fassio, the reason why it would not be rational for Catriona to take the bet has nothing to do with the stakes. Put roughly, the thought is that bet cases raise certain factors to salience, and it is these factors that explain why it would not be rational for Catriona to take the bet. To make this point more persuasive, he considers two cases:

Great Importance (GI): Mark is on a train going to Sterling. It is very important that Mark be at the station soon because he is bringing a medicine to his sister Jane. If Mark cannot meet her at the station and give her the medicine, she is very likely to die. If he doesn't get off the train at right station (e.g., he gets off at a different station), he will not be able to meet his sister on time. According to the route plan the next stop is Sterling. A ticket inspector passed in the coach a few minutes ago and announced that the next stop is Sterling. The train is approaching a station. Mark forms the outright belief that the train is approaching Sterling station, stands up and goes to the exit in order to get out

High Bet (HB): Mark is on a train going to Sterling. As in GI, according to the route plan the next stop is Sterling. A ticket inspector passed in the coach a few minutes ago and announced that the next stop is Sterling. A man approaches him. He proposes a bet about whether the next station is Sterling: if it is, he will give Mark £10000, otherwise Mark will lose £10000. The train is approaching a station. Mark has the same evidence as in GI, but doesn't feel confident enough to take the bet (forthcoming, 14).

The thought is that it would be rational for Mark to believe that the next stop is Sterling, and to act on the basis of this belief, in GI, but it would not be rational for Mark to take the bet in HB. But the stakes are actually *higher* in GI than in HB (and, as Fassio says, if you don't think they are, just make the stakes even higher in GI). So it can't be stakes that are driving intuitions in bet cases.

Now, this does not establish that our perceptual, memorial or testimonial beliefs are ever practically certain. This leads to the second part of Fassio's argument. His basic strategy is to argue that there are examples of perceptual, memorial and testimonial beliefs of which we are practically certain. Take some of the examples considered above: the belief that there are two oranges on the table, formed in normal perceptual conditions; the belief that Rome is the capital of Italy; Catriona's

belief that the next stop is South Parkway. I agree with Fassio that he can accommodate these sorts of cases. But how does this fare as an attempt to satisfy the key desideratum?

The most Fassio can claim is that his account can accommodate some core cases of perceptual, memorial and testimonial knowledge (the cases discussed above). This is not to say that, on his account, we have *as much* perceptual, memorial and testimonial knowledge as we would ordinarily assume. But recall my objection to Kyriacou's radical skeptical invariantism: the more revisionary of our ordinary understanding a view is, the greater its theoretical benefits need to be. Because Kyriacou's view is extremely revisionary of our ordinary understanding of the extent of knowledge, it needs to have far greater theoretical benefits than competitors. Because Fassio's view is far less revisionary, it doesn't need to have as many theoretical benefits. This of course is not to say that it has sufficient theoretical benefits. But the *barrier* to acceptance is considerably lower.

How does Fassio's view fare with introspective beliefs, or beliefs in a priori, analytic or necessary truths? Things are far simpler here. While he doesn't consider such beliefs in much depth, one might think it is very plausible that we can be *absolutely* certain when it comes to beliefs in either of these classes. For Fassio, absolute certainty entails, but is not entailed by, practical certainty. So Fassio's view does well here.

Finally, we can consider the sixth and final class of beliefs I said an account of knowledge should accommodate: certain scientific beliefs, like the belief that climate change is caused by human activities. I am going to argue that it is plausible this belief is practically certain in Fassio's sense, at least for some epistemic agents.

Let's imagine an epistemic agent, Ailsa, who has a keen amateur interest in climate science. For a layperson, she is very well informed about the scientific evidence that human activities are causing global warming. She has read some journal articles, knows all the most common climate skeptical arguments, and can explain, in some detail, where these arguments go wrong. She is also a passionate environmentalist. But her acceptance of the science on climate change isn't dogmatic: she recognises that, while the basic hypothesis that humans are causing global warming is as established as any hypothesis could be, there are still a lot of unanswered questions and uncertainties. I would submit that Ailsa knows that human activities are causing global warming. So, the question is, is her belief is practically certain?

For Ailsa's belief to be practically certain it would have to be that, no matter what the stakes were, if she were rationally responsive to the stakes, she would remain confident enough to believe that human activities are causing global warming, and to act on this basis. Is this the case? At least initially, it might seem like the answer

is clear: yes! After all, you might think, the stakes are incredibly high: the survival of humankind may well depend on dealing with global warming. So, if we think Ailsa is rational in believing that human activities are causing global warming, and acting on this basis, in the scenario described above, then we presumably would think the same no matter the stakes. After all, how could they be higher?

There are three issues. The first issue concerns what the stakes actually are here. Of course, in some sense, the stakes are clearly incredibly high: the survival of humankind may well depend on dealing with global warming. But, in the literature on interest-relative invariantism, when people talk about the stakes relating to a proposition, they generally mean the *costs of being wrong* in believing that proposition.<sup>17</sup> Applying this here, what would matter would be the costs (for Ailsa) of being wrong in believing that human activities are causing global warming. It isn't so clear that the costs (at least for Ailsa) of being wrong in believing this are particularly high.

Perhaps this issue can be dealt with by slightly changing how we think about stakes. Instead of just considering the costs (for a subject) of being wrong in believing some proposition, we can also consider the costs (for that subject) of *not being right*.<sup>18</sup> That is, we can also consider the costs of not believing something that is true. If we put things this way, it may become easier to see why the stakes are high for Ailsa: she presumably would think that the costs of not believing that human activities are causing global warming, when they in fact are doing so, are very high.

The second issue is that, even if we think of the stakes in terms of the costs of not being right, it isn't so clear that the stakes are high for Ailsa here. What bad consequences would follow if Ailsa didn't believe that human activities were causing global warming? Of course, if lots of people don't believe this, that might be bad for society, and detrimental to the long-term survival of the human race. But that's different from it being bad for Ailsa.

To address this issue, we need to further consider the way in which stakes are usually thought of in the literature. We can distinguish between the "perceived" and "actual" stakes for a subject. The perceived stakes are, roughly, what they take the costs of being wrong (and, we might add, of not being right) to be. The actual stakes are, roughly, what the consequences of being wrong will actually be for the subject. Some (e.g. Fassio) interpret stakes in the first way; others (Stanley 2005) in

---

<sup>17</sup> For discussion of the lack of clarity about what stakes are in this literature see Anderson (2015) and Anderson and Hawthorne (2019).

<sup>18</sup> For similar points see Crewe and Jenkins Ichikawa (forthcoming), Mueller (2017), Pace (2011) and Worsnip (2015).

the second. Applying this to Ailsa, it seems clear that the stakes are high in the first sense, but it is less clear that they are high in the second sense.

Does this matter? It is important to note that, even though the notion of stakes may be difficult to apply in Ailsa's case, it is not clear it makes a difference so far as practical certainty is concerned. Imagine you thought that, contrary to first impressions, Ailsa is actually not in a high stakes situation. Does that mean she isn't practically certain? Well, she would only not be practically certain if there were a high stakes situation in which, if she were rationally responsive to the stakes, she would not remain confident enough to believe and act. While I have no argument that there couldn't be such a situation, that is not to say that there actually is one.

The third issue is whether what I have said about Ailsa is going to generalise. It clearly isn't going to generalise in the sense that everyone can be said to know that human activities are causing global warming. But we wouldn't want it to generalise in this way, because some people don't know this. Some people don't know it because they don't believe it. Some people don't know it because, even though they believe it, they believe it on the basis of weak or possibly no evidence.<sup>19</sup> There would only be a problem if there were subjects who believe it on the basis of evidence we would regard as strong enough for knowledge, but aren't practically certain. Are there such subjects? I'm not sure. In any case, if it turns out that, on Fassio's account, some subjects don't know that human activities are causing global warming because of how they would (rationally) respond to the stakes in counterfactual scenarios, then this could be the basis of an objection to Fassio's account of knowledge. But it isn't clear that the objection would have anything to do with the skeptical aspect of his account. The objection would rather be that it just doesn't seem like how you would respond to the stakes in counterfactual scenarios should make a difference to whether you know these things.

#### **4. Concluding Remarks**

This paper started in a rather unpromising way. I noted that it is a desideratum on any account of knowledge that it accommodate the fact—and I think it is a fact—that we know quite a lot of things. But, while this rules out radical forms of skeptical invariantism, it doesn't automatically rule out more moderate forms. I have proposed a list of things I think any account of knowledge should say that we know, and considered whether Fassio's moderate skeptical invariantist account can accommodate them. I have argued that the answer is a qualified "yes". While Fassio focuses on showing that his view can accommodate some perceptual, memorial and testimonial beliefs, I have argued that his account accommodates certain introspective beliefs, beliefs in a priori, analytic or necessary truths and, finally, certain scientific beliefs. This is in no way a "complete" defence of

---

<sup>19</sup> For an empirically informed discussion of this see McKenna (2019).

moderate skeptical invariantism, or Fassio's specific version of it. But I hope I have succeeded in showing that it is plausible to hold that, even though we know a fair bit, we know quite a lot less than we ordinarily take ourselves to know.<sup>20</sup>

## References

- Anderson, Charity. 2015. "On the Intimate Relationship of Knowledge and Action." *Episteme* 12 (3): 343–53.
- Anderson, Charity, and John Hawthorne. 2019. "Knowledge, Practical Adequacy, and Stakes." *Oxford Studies in Epistemology* 6: 234–57.
- Brown, Jessica. 2006. "Contextualism and Warranted Assertibility Manoeuvres." *Philosophical Studies* 130 (3): 407–35.
- . 2008. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Noûs* 42 (2): 167–89.
- . 2018. *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press.
- Cappelen, Herman. 2005. "Pluralistic Skepticism: Advertisement for Speech Act Pluralism." *Philosophical Perspectives* 19 (1): 15–39.
- Cohen, Stewart. 1986. "Knowledge and Context." *Journal of Philosophy* 83 (10): 574–83.
- Cook, John, Naomi Oreskes, Peter T. Doran, William R. L. Anderegg, Bart Verheggen, Edward H. Maibach, J Stuart Carlton, Stephan Lewandowsky, Andrew G. Skuce, and Sarah A. Green. 2016. "Consensus on Consensus: A Synthesis of Consensus Estimates on Human-Caused Global Warming." *Environmental Research Letters* 11 (4): 048002.
- Craig, Edward. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- Crewe, Bianca, and Jonathan Jenkins Ichikawa. forthcoming. "Rape Culture and Epistemology." In *Applied Epistemology*, edited by Jennifer Lackey. Oxford: Oxford University Press.
- Davis, Wayne A. 2007. "Knowledge Claims and Context: Loose Use." *Philosophical Studies* 132 (3): 395–438.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52 (4): 913–29.
- . 2017. *The Appearance of Ignorance: Knowledge, Skepticism, and Context, Vol. 2*. Oxford: Oxford University Press.
- Dinges, Alexander. 2016. "Skeptical Pragmatic Invariantism: Good, but Not Good Enough." *Synthese* 193 (8): 2577–93.
- Fantl, Jeremy, and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- Fassio, Davide. forthcoming. "Moderate Skeptical Invariantism." *Erkenntnis*, 1–30.

---

<sup>20</sup> Thanks to Davide Fassio, Michael Hannon and Christos Kyriacou for comments on an earlier version of this paper.

- Fassio, Davide, and Robin McKenna. 2015. "Revisionary Epistemology." *Inquiry: An Interdisciplinary Journal of Philosophy* 58 (7–8): 755–79.
- Fumerton, Richard A. 1995. *Metaepistemology and Skepticism*. Rowman & Littlefield.
- Gerken, Mikkel. 2017. *On Folk Epistemology. How We Think and Talk about Knowledge*. Oxford University Press.
- Hannon, Michael. 2020. "Why Purists Should Be Infallibilists." *Philosophical Studies* 177 (3): 689–704.
- Haslanger, Sally. 1999. "What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology." *Philosophical Perspectives* 13 (s13): 459–80.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hazlett, Allan. 2009. "Knowledge and Conversation." *Philosophy and Phenomenological Research* 78 (3): 591–620.
- Kusch, Martin, and Robin McKenna. forthcoming. "The Genealogical Method in Epistemology." *Synthese*.
- Kyriacou, Christos. 2017. "Bifurcated Sceptical Invariantism." *Journal of Philosophical Research* 42: 27–44.
- Levin, Janet. 2008. "Assertion, Practical Reason, and Pragmatic Theories of Knowledge." *Philosophy and Phenomenological Research* 76 (2): 359–384.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–67.
- McKenna, Robin. 2019. "Irrelevant Cultural Influences on Belief." *Journal of Applied Philosophy* 36 (5): 755–68.
- Moore, G. E. 1939. "Proof of an External World." *Proceedings of the British Academy* 25: 273–300.
- Mueller, Andy. 2017. "Pragmatic or Pascalian Encroachment?" *Logos and Episteme* 8 (2): 235–41.
- Pace, Michael. 2011. "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe.'" *Noûs* 45 (2): 239–68.
- Pritchard, Duncan. 2002. "Resurrecting the Moorean Response to the Sceptic." *International Journal of Philosophical Studies* 10 (3): 283 – 307.
- Quine, W. V. O. 1969. "Epistemology Naturalized." In *Ontological Relativity and Other Essays*, 69–90. New York: Columbia University Press.
- Reed, Baron. 2010. "A Defense of Stable Invariantism." *Noûs* 44 (2): 224–44.
- Rysiew, Patrick. 2001. "The Context-Sensitivity of Knowledge Attributions." *Noûs* 35 (4): 477–514.
- Sosa, Ernest. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13 (s13): 137–49.
- . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- . 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford University Press.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.

- Stroud, Barry. 1984. *The Significance of Philosophical Scepticism*. Oxford University Press.
- Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. Oxford University Press.
- Weatherson, Brian. 2003. "What Good Are Counterexamples?" *Philosophical Studies* 115 (1): 1–31.
- . 2012. "Knowledge, Bets, and Interests." In *Knowledge Ascriptions*, edited by Jessica Brown and Mikkel Gerken, 75–103. Oxford University Press.
- Williams, Michael. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Princeton University Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Worsnip, Alex. 2015. "Two Kinds of Stakes." *Pacific Philosophical Quarterly* 96 (3): 307–24.
- Wright, Crispin. 1991. "Scepticism and Dreaming: Imploding the Demon." *Mind* 100 (1): 87–116.