

AGAINST MIXED EPISTEMOLOGY

JOE MILBURN

Abstract. We can call any reductive account of knowledge that appeals to both safety and ability conditions a mixed account of knowledge. Examples of mixed accounts of knowledge include Pritchard's (2012) *Anti-Luck Virtue Epistemology*, Kelp's (2013) *Safe-Apt account of knowledge*, and Turri's (2011) *Ample belief* account of knowledge. Mixed accounts of knowledge are motivated by well-known counterexamples to pure safety and pure ability accounts of knowledge. It is thought that by combining both safety and ability conditions we can give an extensionally adequate reductive account of knowledge. In this paper I argue that the putative counterexamples to pure safety and pure ability accounts of knowledge fail to motivate mixed accounts of knowledge. In particular, I argue that if the putative counterexamples are problematic for safety accounts they are problematic for ability accounts and vice-versa. The reason for this, I argue, is that the safety condition and ability condition should be understood as alternative expressions of the same intuition — that knowledge must come from a reliable source.

Keywords: Anti-Luck epistemology; virtue epistemology; reliabilism.

1. Introduction

Many philosophers believe that there is a *safety condition* on knowledge. The exact formulation of the safety condition is controversial, but the basic idea is that if S knows that *p*, then S forms her belief that *p* in such a way that she is not at risk of forming a false belief (for this formulation of the core idea of safety, cf. Bogardus 2014).

There are also many philosophers who believe there is an *ability condition* on knowledge. Again, the exact formulation of the ability condition is controversial, but the basic idea is that if S knows that *p*, then S's true belief is creditable to a cognitive ability of S's (cf. Sosa 2007; Schafer 2014).

Philosophers have tried to give a reductive account of knowledge solely in terms of the safety condition (e.g. Pritchard 2005; 2007) or solely in terms of the ability condition (e.g. Sosa 2007; Greco 2010; 2012, and Turri 2011; 2015). Following Pritchard (2012) we can respectively call these different reductive accounts *pure anti-luck accounts* and *pure virtue accounts*. In turn we can call any reductive account of knowledge that appeals to both safety conditions and ability conditions a *mixed account of knowledge*.

Recently, Duncan Pritchard (2012) has argued for a mixed account of knowledge, defending what he calls *Anti-Luck Virtue Epistemology* (ALVE). According to ALVE:

Principia 19(2): 183–195 (2015).

Published by NEL — Epistemology and Logic Research Group, Federal University of Santa Catarina (UFSC), Brazil.

S knows that *p* if and only if S's safe true belief that *p* is the product of her relevant cognitive abilities (such that her safe cognitive success is to a significant degree creditable to her cognitive agency) (Pritchard 2012, p.273).

Christoph Kelp (2013) defends his own mixed account of knowledge, which he calls the *Safe-Apt* view of knowledge. Says Kelp:

The core thesis of virtue epistemology is that knowledge involves cognitive success that is due to cognitive competence. More specifically ... One knows that *p* only if one believes *p* and the acquisition and/or retention of one's true belief that *p* is (sufficiently) due to cognitive competence or ability ... I want to suggest that combining ... an additional safety condition with a Sosa-style competence-manifestation account of the 'because' relation [in the core thesis of virtue epistemology] can avoid a number of problems that beset its competitors (Kelp 2013, p.266).

John Turri (2011) has considered another mixed account of knowledge. According to Turri, knowledge might be *ample* belief; i.e safe belief in which the safety of the belief manifests the believer's cognitive ability.¹

There are two reasons that philosophers interested in reductive accounts of knowledge have been attracted to mixed accounts of knowledge. First, both pure safety accounts and pure ability accounts struggle with well-known counterexamples. However, the counterexamples that plague pure safety accounts do not seem to affect pure ability accounts, and likewise, the counterexamples that plague pure ability accounts do not seem to affect pure safety accounts. The thought is that by combining both safety and ability conditions we can give an extensionally adequate reductive account of knowledge. Second, philosophers like Pritchard and Kelp argue that given the function of knowledge ascriptions, we should expect that there would be distinct safety and ability conditions on knowledge.

In this paper I want to focus on the first motivation for mixed accounts of knowledge, i.e., that while pure safety and pure ability accounts struggle to achieve extensional adequacy, mixed epistemology does not. I argue that adding an ability condition to the safety condition and *vice versa* does nothing to help with the putatively problematic cases. This is because both the safety condition and the ability condition are ways of spelling out the intuition that knowledge must come from reliable methods of belief formation in the right way; furthermore, the putative counterexamples to both pure safety accounts and pure ability accounts arise because of problems inherent in our notion of reliability.

The paper proceeds as follows. In the first section I will look at the putative counterexamples for pure safety and pure ability accounts that are used to motivate mixed accounts of knowledge. In the second section I will argue that these sorts of cases are equally problematic (or unproblematic) for both pure safety and pure ability ac-

counts. In the third section, I will argue that the putative counterexamples arise because both the safety and ability conditions are attempts to spell out the intuition that knowledge must come from a reliable source, and attributions of reliability are context sensitive.

2. The Cases

Mixed accounts of knowledge are motivated by counterexamples to pure safety accounts of knowledge which are not counterexamples to pure ability accounts of knowledge and *vice versa*. What are these counterexamples?

First there are so-called dangerous Gettier cases. The most famous of these cases is the Red Barn Façade case. Here is Pritchard's formulation.

Barney

Using his reliable perceptual faculties, Barney non-inferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact barn facades (Pritchard 2012, p.251).

Pure virtue epistemologies are supposed to struggle with ruling out dangerous Gettier cases like the *Barney case*. In this case, Barney is said to form the true belief that there is a red barn in front of him using his reliable perceptual faculties; we can further stipulate that lighting conditions etc. are normal. It follows that Barney's true belief that there is a red barn in front of him is attributable to a reliable method of belief formation; furthermore, this reliable method of belief formation is a cognitive ability. After all, reliable perceptual faculties are paradigmatic instances of epistemic abilities. Thus, *prima facie*, in the *Barney case* and other similar cases, pure ability accounts must rule that Barney has knowledge.²

Epistemologists defending a mixed account of knowledge assume that, while the ability condition struggles to rule out dangerous Gettier cases, the safety condition easily does so. After all, Barney is at risk for forming a false belief — he could have easily been looking at a barn façade and in this case he would have formed a false belief that there was a barn in front of him.

If the safety condition deals easily with dangerous Gettier cases like the *Barney case*, however, it struggles to deal with other sorts of cases in which one's true belief is safe, but safe for the wrong sorts of reasons. Consider for example, Pritchard's *Temp case*.

Temp

Temp forms his beliefs about the temperature in the room by consulting a

thermometer. His beliefs, so formed, are highly reliable, in that any belief he forms on this basis will always be correct. Moreover, he has no reason for thinking that there is anything amiss with his thermometer. But the thermometer is in fact broken and is fluctuating randomly within a given range. Unbeknownst to Temp, there is an agent hidden in the room who is in control of the thermostat whose job it is to ensure that every time Temp consults the thermometer the ‘reading’ on the thermometer corresponds to the temperature in the room (Pritchard 2012, p.260).

Prima facie, Temp’s belief is safe, but it isn’t an instance of knowledge. So pure safety accounts of knowledge are forced to give the wrong verdict regarding the *Temp case*.

Epistemologists giving a mixed account of knowledge assume that the ability condition easily rules out cases like the *Temp case*. The thought is that in the *Temp case*, Temp’s belief is not creditable to any method of belief formation that can be called a virtue or an ability. As Pritchard puts it, “[the] correctness [of Temp’s beliefs] has nothing to do with Temp’s abilities and everything to do with some feature external to his cognitive agency” (Pritchard 2012, p.260).

So while the ability condition struggles to rule out cases like the *Barney case*, it does a fine job with the *Temp case*, and while the safety condition struggles with the *Temp case*, it does a fine job with the *Barney case*. Mixed accounts of knowledge take advantage of these putative facts and combine both safety and ability conditions in a reductive account of knowledge.

3. Why the Cases Don’t Motivate Mixed Accounts of Knowledge

In the previous section I reviewed some cases that are supposed to motivate mixed accounts of knowledge. In this section, I argue that, in fact, these cases do *not* motivate mixed accounts of knowledge. If ability conditions on knowledge struggle to rule out dangerous Gettier cases like the *Barney case*, so do safety conditions; if safety conditions on knowledge struggle to rule out cases like the *Temp case*, so do ability conditions.

First, I will consider dangerous Gettier cases like the *Barney case*. Despite initial appearances to the contrary, the safety condition on knowledge struggles to rule out dangerous Gettier cases as instances of knowledge. To see this, consider the following case from Comesaña (2005) that he brings up as a counterexample to the safety condition on knowledge.

HALLOWEEN PARTY: There is a Halloween party at Andy’s house, and I am invited . . . [Andy] hires Judy to stand at a crossroads and direct people towards the house . . . Unbeknownst to me, Andy doesn’t want Michael to go to the party, so he also tells Judy that if she see Michael she should tell him the same thing she tells everybody else (that the party is at the house down the

left road), but she should immediately phone Andy so that the party can be moved to Adam's house, which is down the right road. I seriously consider disguising myself as Michael, but at the last moment I don't. When I get to the crossroads, I ask Judy where the party is, and she tells me that it is down the left road (Comesaña 2005, p.397).

In the Halloween case, you have knowledge, but your belief doesn't seem to be safe: i.e., you seem to be at risk of forming a false belief. So it seems that there is no safety condition on knowledge.

The right response to this kind of counterexample is to *relativize* the safety condition to a method of belief formation. In this case beliefs are not safe *simpliciter*, but are safe *relative* to the method by which they were formed (cf. Pritchard 2007; 2009). Given that you form your belief as to where the party is on the basis of *knowledgeable and sincere testimony*, you are not at risk of forming a false belief. So *pace* Comesaña your belief in the *Halloween party case* is an instance of knowledge, and it is safe. The *Halloween party case* only appears to threaten the safety condition on knowledge because you were at risk of not forming your belief on the basis of *knowledgeable and sincere testimony*. If, however, we don't confuse being at risk of forming a false belief given one's method of belief formation, and being at risk of not using a safe method of belief formation, then counterexamples like the *Halloween party case* pose no threat to the safety condition on knowledge (cf. Bogardus 2014, p.297–9).

However, when we relativize the safety condition to a method of belief formation, using the safety condition to rule out Red Barn facades and other dangerous Gettier cases becomes problematic. To see this, consider a method *m* such that using *m* entails the truth of *p*. We can call such methods *factive* methods. Properly deducing *p* from *q* when one knows that *q* is an example of such a method. Other factive methods include believing that *p* on the basis of seeing that *p*, believing that *p* on the basis of sincere, competent, and knowledgeable testimony that *p*, and believing that *p* on the basis of correctly remembering that *p*. Every method of belief formation that is a factive method produces *only* maximally safe beliefs. Given that you form your belief on the basis of seeing that *p*, you are at absolutely no risk of forming a false belief that *p* since, given that you form your belief on the basis of seeing that *p*, your belief that *p* must be true.

We can give a plausible description of the *Barney case* in which Barney forms his belief that there is a red barn in front of him on the basis of *seeing that* there is a red barn in front of him. Under this description, Barney's belief will be safe. For again, given that Barney's belief that there is a red barn in front of him is based on his seeing that there is a red barn in front of him, there is *no* risk that Barney will believe falsely.

If safety is to rule out the *Barney case* then it must be assumed that the method that Barney is using is non-factive. Suppose for the sake of argument that this assumption is justified. Then it seems that the safety condition implies that Barney

lacks knowledge. We can, for the sake of illustration, suppose that Barney bases his belief that there is a red barn in front of him on the fact that he has no reason for doubting that he is not in normal conditions, and it appears to him that there is a red barn in front of him. (The exact details of the method don't matter for our purposes). In this case, it seems that Barney's belief that there is a red barn in front of him is not safe. In many nearby worlds, worlds in which he just happens to be facing one of the many barn facades littering the countryside, Barney forms false beliefs using the same method, and so Barney is at risk of forming a false belief given the method of belief formation he is using.

But even on the assumption that Barney's method of belief formation is non-factive, there are problems. Consider the following case.

Red Barn County Architect: Peter is an epistemic thief. At great personal expense, and after years of research and development, he created the Red Barn Façade County Erector Set. Once the erector set was put in place, at the push of button, any rural county could be turned into the epistemic equivalent of Red Barn Façade County. (The erector set was amazing — it could instantly raze all but one of the red barns in the county and set up any number of red barn facades in their place.) One day, it turns out that Barney is driving through a rural county in which Peter had set up his erector set. Of course Barney is unaware of this. He is about to fall into the trap. With Barney approaching a red barn that is to be razed and replaced by a red barn facade, Peter pushes the button! Then the completely unexpected happens. Peter's almost perfectly reliable erector set fails to work. Barney sees the red barn in front of him and forms the belief: there is a red barn in front of me.

Here it seems that Barney obviously has knowledge that there is a red barn in front of him. But in many of the nearest possible worlds in which Barney forms his belief that there is a red barn in front of him on the basis that has no reason for believing that conditions are not normal and it seems to him that there is a red barn in front of him, Barney forms a false belief. So it can seem that in the *Red Barn County Architect case* Barney is at risk of forming a false belief if we assume his method of belief formation is not factive.

These reflections show that the safety condition on knowledge cannot deal with dangerous Gettier cases in a simple way. The problem is that we cannot say whether a belief is safe until we have fixed the relevant method of belief formation, but in dangerous Gettier cases (e.g. the *Barney case*) and in cases that are almost dangerous Gettier cases (e.g. the *Red Barn Façade Architect case*) we can give alternative descriptions of the relevant methods of belief formation that make the belief respectively safe or unsafe, where these alternative descriptions are equally plausible.

In effect, we have an analogue to the generality problem for process reliabilism. While various solutions to the generality problem have been put forward, the one

that comes most easily to hand is Heller's (1995) *contextualist* or *simple* response. According to Heller there is no fact of the matter whether a given belief is safe or not, as there is no fact of the matter as to what methods of belief formation are relevant for the epistemic assessment of a belief.³ Rather, there are contextual factors that determine what the relevant methods of belief formation are for deciding whether a belief is safe. In this case, it won't be true that a belief is formed on the basis of some method *m* absolutely speaking but only relative to some context of knowledge attribution.⁴

Of course, the details of this contextualist treatment of the safety condition are controversial and tough to make out. The important point is that the kind of contextualist maneuver needed by safety theorists to successfully deal with the *Barney case* could also allow virtue theorists to give the desired verdict. This is especially clear if we think of the methods of belief formation relevant to credit theorists as being reliable methods of belief formation relative to some environment. Indeed Greco (2010), (2012) appeals to contextualist ideas to allow his pure virtue account to deal with dangerous Gettier cases. Once again, the details of this contextualist treatment of epistemic ability are controversial, but in this Greco's pure virtue account of knowledge is no worse off than pure safety accounts of knowledge. The upshot of these reflections, then, is that adding safety conditions to pure virtue accounts of knowledge gives them no real advantage in ruling out cases like the *Barney case*.⁵

Of course, mixed epistemologists might reject Heller's contextualist solution to the generality problem. But in this case, they must provide their own solution to the problem. Until this is done, it is unclear how adding a safety condition to knowledge helps to rule out dangerous Gettier cases.

Next consider cases in which one's belief appears to be safe, but is safe for the wrong reasons, such as the *Temp case*. Despite initial appearances to the contrary, the ability condition on knowledge fails to rule out the *Temp case* as a case of knowledge in a straightforward manner.

Pritchard believes that *Temp* fails to meet the ability condition on knowledge, because "[the] correctness [of *Temp*'s beliefs] has nothing to do with *Temp*'s abilities and everything to do with some feature external to his cognitive agency" (Pritchard 2012a, p.260). But this is not true. The correctness of *Temp*'s beliefs does seem to be attributable, at least in part, to *Temp*'s cognitive agency. After all, *Temp* wouldn't form the correct belief if he didn't have the ability to read the thermostat, or if he didn't think that it was rational to trust the thermostat, and instead thought the rational thing to do was to believe that the temperature in the room is anything but what the thermostat read. And it seems that if *Temp* was so mad as to form any random belief on the basis of what the thermostat reads, e.g. if he was disposed to form beliefs regarding the price of oil futures on the basis of reading the thermostat, we couldn't so much as consider his belief to be safe. So *Temp*'s forming the correct

beliefs regarding the temperature in the room is attributable, in some part at least, to his cognitive agency.

If appeals to the ability condition are to rule out the *Temp case*, they must do so for reasons other than the putative fact that Temp is in no way creditable for his true belief. But it seems that this cannot be done by demanding that *the greater* share of credit for one's cognitive success belong to the knowledgeable subject. As Lackey (2009) and others have noted, this move seems to imply that we cannot have knowledge on the basis of testimony from experts, since in these cases the greater share of credit for one's cognitive success goes to the expert and not to the one who accepts the testimony.

Another possibility that has been explored in the literature is that we should understand the ability condition in terms of 'manifestation'. On this view, believing truly that *p* can be creditable to an agent *S*, just in case in believing truly that *p*, *S* manifests her cognitive ability (cf. Turri 2011; Sosa 2010). This is exactly the line of thought that Kelp (2013) picks up on. On this way of thinking, Temp's true beliefs regarding the temperature in the room fail to be an instance of knowledge, because Temp's cognitive success in these instances fails to manifest his cognitive ability. We can contrast this with Temp's beliefs about the temperature formed on the basis of reading the thermostat in normal conditions. In this case, Temp's cognitive success *does* manifest his cognitive ability.

Prima facie this manner of dealing with the problem shows some promise, but I think it ultimately fails. Consider that Temp's believing truly in the *Temp case* manifests a disposition of his to get things right regarding room temperature in conditions described by the *Temp case*. If the ability condition is to rule out that Temp has knowledge, then it needs to be able to rule out the possibility that this disposition is a cognitive ability of Temp's.

A possible way of dealing with this problem is to hold that epistemic abilities must be an integrated part of one's cognitive character. Greco (2010) takes this tack in dealing with the so-called problem of strange and fleeting processes, and Pritchard follows suit. Says Pritchard:

We ... need to ask what it is about the particular belief-forming dispositions that qualify as cognitive abilities that makes them knowledge conducive. I take it that as a minimal requirement these belief-forming dispositions should be both reliable and suitably integrated with the agent's other belief-forming dispositions (Pritchard 2012, p.262).

But what is it for a belief-forming disposition to be suitably integrated with the agent's other belief forming dispositions? And why shouldn't we consider the relevant dispositions to get things right in the *Temp case* to be well-integrated belief-forming dispositions?

According to Greco cognitive integration is a matter of both a range of outputs and how this range of outputs interacts with other belief forming dispositions, as well as sensitivity to defeating evidence. Says Greco:

One aspect of cognitive integration concerns the range of outputs — if the product of a disposition are few and far between, and if they have little relation to other beliefs in the system, then the disposition is less well integrated on that account. Another aspect of cognitive integration is sensitivity to counterevidence, or to defeating evidence. If the beliefs in question are insensitive to reasons that count against them, then this too speaks against cognitive integration (Greco 2010, p.152).

Arguably, Temp's disposition to get things right in situations like those described in the *Temp case* is sufficiently integrated into Temp's cognitive character to count as an ability. On the one hand, Temp's cognitive disposition has a fairly wide range of outputs - it is a disposition to form a wide range of beliefs concerning the temperature in the room and these beliefs concerning room temperature might interact with any number of beliefs formed by Temp's other epistemic dispositions. On the other hand, there is no reason to suppose that the beliefs generated by Temp's disposition are insensitive to counterevidence. We can suppose that Temp might disbelieve that the room temperature is what the thermostat says it is if he has evidence that the thermostat is broken. Given this, it is plausible that Temp's disposition to get things right in cases like the *Temp case* is part of Temp's cognitive character. Temp is the sort of person who will get things right in these sorts of situations; we can contrast Temp again with someone who forms absolutely irrational beliefs on the basis of reading a thermostat — whereas it is part of Temp's cognitive character to form the correct belief regarding room temperature in cases like the *Temp case*, it is not part of this subject's cognitive character to do so.

It seems, then, that there is no easy way for ability conditions to rule out the *Temp case* as an instance of knowledge. But this should come as no surprise, for well-known reasons. The ability condition's struggle to successfully rule out knowledge in the *Temp case* is just another manifestation of the generality problem for reliabilist epistemologies. Temp's token belief can be said to be the manifestation of many different dispositions, some of which will be reliable, others which will not be. The difficulty is giving a principled reason for saying why a certain disposition *is* or *is not* relevant. Pure virtue accounts of knowledge attempt to solve the problem by claiming that the relevant dispositions are *abilities*. But absent an account of cognitive abilities this only pushes the problem back. And as the remarks above show, it is difficult to give an account of cognitive ability that could be put to use in deciding cases like the *Temp case*. A more promising tack is to give a contextualist account of cognitive ability, and this is exactly what Greco does in more recent work. Says Greco:

Knowledge-relevant abilities have the following structure.

S has a knowledge-relevant ability $A(R/C/D)$ relative to an environment $E = S$ has a disposition to believe truths in range R when in circumstances C and environment E , with degree of reliability D .

How do we fill in the relevant parameter? For example, what counts as the relevant range of truths, or the relevant sort of circumstances? [...] I want to propose [...] that we look to relevant information needs to set our parameters. That is, the parameters are set by the needs for information and information sharing attaching to relevant practical tasks (Greco 2012, p.17–8).

Greco then goes on to note that the relevant information needs can be either those of the epistemic subject or those who would attribute knowledge to the epistemic subject (Greco 2012, p.19). Greco's account of knowledge-relevant ability, then, is broadly contextualist: there is no simple fact of the matter whether some cognitive disposition is a knowledge-relevant ability; rather a given disposition will be a knowledge-relevant ability relative to the informational needs of epistemic subjects and epistemic attributors.

While this broad contextualism allows for the ability condition to successfully rule that Temp fails to have knowledge, a similar kind of account that replaces talk of *knowledge-relevant abilities* with *safety* can also rule out that Temp has knowledge. For given our informational needs it is natural to describe Temp's method of belief formation as forming a belief on the basis of a broken measuring instrument, in which case it follows that he is at serious risk of forming a false belief. And again, given our informational needs we should treat possible worlds in which there is no one manipulating the temperature to match the thermostat as nearby when evaluating the *Temp case*, in which case it follows again (on plausible assumptions) that Temp is at serious risk of forming a false belief. So while a broadly contextualist construal of the ability condition can rule out the Temp case as an instance of knowledge, so too can a broadly contextualist construal of the safety condition.

Of course, the mixed epistemologist might reject Greco's contextualism for knowledge-relevant abilities. But in this case, she must provide some other account of what makes a cognitive disposition an ability. Until she does so, both safety and virtue theories are equal footing when it comes to ruling out cases like the *Temp case*. The upshot of this is that if cases like the *Temp case* are a real problem for safety conditions to deal with, appealing to ability conditions does nothing to help.

4. Diagnosis and Conclusion

In the past section I have argued that cases like the *Barney case* and cases like the *Temp case* are equally problematic (or unproblematic) for pure safety and pure virtue

accounts of knowledge. Why is this so? The answer seems to be that both safety conditions and ability conditions are different ways of spelling out the intuition that knowledge must come from a reliable source in the right way (cf. Goldman 2012, p.70, p.87–90). But if contextualists like Heller (1995) are right, there is no fact of the matter whether or not a given belief is reliable. Relative to one context of evaluation, Barney's belief is highly reliable; from another context of evaluation, it is not. Again, relative to one context of evaluation, Temp's belief is highly reliable, relative to a different context of evaluation it is not. It seems that it is the relativity of reliability that allows the putative counterexamples to safety conditions and ability conditions to get a grip. The counterexamples are compelling only to the extent that we fail to notice that we shift from one context of evaluation to another when we judge whether a belief is safe, or formed by an ability, and when we judge whether the belief is an instance of knowledge. But if we keep the context of evaluation fixed by starting out with a judgment that an individual lacks knowledge in a given case, we can go on to describe the individual as failing both to meet the ability and safety condition. And likewise, if we keep the context of evaluation fixed by starting out with a judgment that an individual possesses knowledge in a given case, we can go on to describe the individual as meeting both the ability and safety conditions on knowledge. My point is that if the counterexamples to both pure safety and pure ability accounts of luck arise because of the context sensitivity of reliability attributions merely combining safety and ability conditions will be no help in dealing with these counterexamples.

In this paper I have argued against mixed accounts of knowledge. The main motivation for mixed accounts is that they are supposed to easily deal with cases problematic for their pure rivals. If I am right, this motivation is illusory. Indeed, it must be if both safety conditions and ability conditions are ways of expressing reliability conditions on knowledge.

One possible response to my argument is to reject both safety accounts, virtue accounts and mixed accounts of knowledge, and seek some other way of dealing with these problematic cases. Another possibility is to question whether these cases are problematic after all. There is some reason for pursuing this line; in recent years more and more epistemologists have expressed doubts as to whether knowledge is lacking in so-called dangerous Gettier cases; and I have given reason for supposing that both pure safety accounts and pure ability accounts can both deal with cases like the *Temp* case above. Given this, we might take a no-priority view between safety accounts and ability accounts of knowledge. On this view, meeting the safety conditions on knowledge entails that one meets the ability conditions on knowledge and vice-versa. As a result neither safety nor ability conditions are seen as having priority over the other, but both are seen as different ways of spelling out a fundamental intuition concerning knowledge.

References

- Bogardus, T. 2014. Knowledge Under Threat. *Philosophy and Phenomenological Research* **88**: 289–313.
- Comesaña, J. 2005. Unsafe Knowledge. *Synthese* **146**: 395–404.
- Goldman, A. 2012. Reliabilism. In: *Reliabilism and Contemporary Epistemology*, p.68–94. Oxford: Oxford University Press.
- Greco, J. 2012. A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research* **85**: 1–26.
- . 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Heller, M. 1995. The Simple Solution to the Generality Problem. *Nous* **29**: 501–15.
- Kelp, C. 2013. Knowledge: The Safe Apt View. *The Australasian Journal of Philosophy* **91**: 265–78.
- Lackey, J. 2009. Knowledge and Credit. *Philosophical Studies* **142**: 27–42.
- Lycan, W. 2006. On the Gettier Problem Problem. In: S. Hetherington (ed.) *Epistemology Futures*, p.148–68. Oxford: Oxford University Press.
- Millikan, R. 1995. Naturalist Reflections on Knowledge. In: *White Queen Psychology and Other Essays for Alice*, p.241–64. Cambridge: The MIT Press.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
- . 2007. Anti-Luck Epistemology. *Synthese* **158**: 277–97.
- . 2009. Safety-Based Epistemology: Whither Now? *The Journal of Philosophical Research*: 33–45.
- . 2012. Anti-Luck Virtue Epistemology. *The Journal of Philosophy* **109**: 247–79.
- Schafer, K. 2014. Knowledge and Two-Forms of Accidental Truth. *Philosophy and Phenomenological Research* **89**: 373–393.
- Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Vol. 1*. Oxford: Oxford University Press.
- . 2010. How Competence Matters in Epistemology. *Philosophical Perspectives* **24**: 465–75.
- Szabo-Gendler, T.; Hawthorne, J. 2005. The Real Guide to Fake Barns: A Catalogue of Gifts for your Epistemic Enemies. *Philosophical Studies* **124**: 331–52.
- Turri, J. 2011. Manifest Failure: The Gettier Problem Solved. *The Philosophers Imprint* **8**: 1–11.
- . 2015. From Virtue Epistemology to Abilism: Theoretical and Empirical Developments. In: C. Miller (ed.) *The Character Project*. Oxford: Oxford University Press.

JOE MILBURN
 The University of Pittsburgh
 jcm55@pitt.edu

Notes

¹ Turri has since disavowed this possibility (cf. Turri 2015). It is worth noting that Turri now accepts unequivocally that we have knowledge in Red Barn façade cases.

² How much of a strike this is against such theories is difficult to assess, because many philosophers, and not just pure virtue epistemologists, now attribute knowledge to subjects in dangerous Gettier cases (cf. Sosa 2007; Turri 2011; Lycan 2006; Millikan 1995 *inter alia* for philosophers who believe that subjects have knowledge in these case, cf. Szabo and Hawthorne (2005) for some of the difficulties in denying knowledge to those in dangerous Gettier case).

³ It should be noted that Heller does not explicitly speak of ‘safety’ but instead of ‘reliability’. *Mutatis mutandis*, his remarks apply equally well to safety.

⁴ My appeals to context here are supposed to be neutral between semantic contextualism and subject sensitive invariantism. The point is simply that whether it is true that an individual S forms her belief that *p* using a method *m* depends on factors concerning the attributors circumstances, values, etc.

⁵ If my argument is correct, then requiring that one’s belief be safe does not help one to rule out dangerous Gettier cases like the *Barney* case. This weighs against both weaker and stronger versions of mixed epistemology. Remember that on the weaker version of mixed epistemology, one’s belief must meet the standard virtue/credit conditions for knowledge, and one’s belief must meet the safety condition. On the stronger version of mixed epistemology one’s belief must meet the safety condition, and its doing so must be attributable to one’s epistemic virtue/ability. Of course, the problem is that much more difficult for stronger versions of mixed epistemology. For it is clear that, even if one’s belief is safe in the *Red Barn Façade Architect* case, its safety does not seem to be attributable to one’s epistemic virtue or cognitive ability.