For *Illuminating Errors: New Essays on Knowledge from Non-Knowledge*,
(eds.) R. Borges & I. Schnee, (London: Routledge, *forthcoming*).

KNOWLEDGE FROM ERROR AND ANTI-RISK VIRTUE EPISTEMOLOGY

DUNCAN PRITCHARD

University of California, Irvine

dhpritch@uci.edu

ABSTRACT. There are a class of examples that appear to show that one can gain knowledge via reasoning from a falsehood. The problem that such examples pose is to explain why knowledge can be acquired in such cases given that in general drawing inferences from false beliefs, even if the inferred belief is true, does not result in knowledge. A safety-based rationale is offered in this regard, along with the kind of motivation for safety offered by *anti-risk epistemology*. It is maintained that such a rationale is at best incomplete, however, and that for a fuller account of how knowledge can be acquired from error we need to turn to the theory of knowledge offered by *anti-risk virtue epistemology*.

1. Generally speaking at least, one cannot acquire knowledge by inference from a false belief, even if one happens via this route to form a true belief in the inferred proposition, and even if one's beliefs in the entailing and entailed proposition are both justified. Edmund Gettier's (1963) original cases illustrate this point nicely, as they both essentially concern a subject inferring a true belief from a false, but justified, belief, and yet no-one is inclined to think that the inferred belief, even while justified, amounts to knowledge. To take one of Gettier's own examples, imagine that our hero justifiably believes the false proposition that Jones owns a Ford, and infers from this that either Jones owns a Ford or Smith is in Barcelona. The inferred disjunctive belief is, plausibly, justified too, and we can also stipulate that it happens to be true—on account of the incidental truth of the

second disjunct—but it hardly seems to qualify as knowledge. In particular, it looks to be just a matter of luck that the inferred belief is true, given how it was formed, in that it is just a coincidence that the second disjunct of the inferred belief happens to be true. Epistemic luck of this sort is, however, usually thought to be incompatible with knowledge.

Or consider the following Gettier-style scenario:

> *Detective*
> Chandler, a private investigator, believes that the murderer is Jones, since all the evidence is pointing that way. As a result he infers that the murderer has blood type O, since that is Jones' blood type. Chandler is wrong about the identity of the murderer, but his inferred belief is true regardless, as the murderer does have blood type O.

Chandler doesn't know what he infers, even though what he infers is true and even though he inferred it from a justified belief (such that it is plausible that this inferred belief is also justified). In particular, he doesn't know what he infers because, as in the previous Gettier case, it seems to be just a matter of luck that his belief is true, given how it was formed—i.e., it is just a coincidence that the murderer and the person who Chandler believes is the murderer happen to have the same blood group.

Of course, Chandler could subsequently acquire independent reasons for holding the inferred belief, and hence come to know it, but in such a case the inferred belief would no longer be epistemically based on the inference from a false belief (but based on the new independent grounds instead). The upshot, it seems, is that one can't gain knowledge by inferring it from a falsehood, at least to the extent that the epistemic basis for the inferred belief remains the inference from a falsehood.

There do, however, seem to be some exceptions to this rule. Consider the following, now familiar, case:

> *Handout*
> Ted needs to determine how many handouts he will need for his talk. Carefully counting the number of people present he forms the belief that there are 53 people in the room. He accordingly infers that the 100 handout copies that he has will be sufficient. But Ted miscounted, and there are in fact 52 people in attendance.[1]

Structurally Handout looks very similar to a Gettier-style case like Detective. Ted's initial belief that there are 53 people at the talk is justified, but false, and the inferred belief is also justified, but true. The difference, however, is that a number of epistemologists have argued that the inferred belief amounts to knowledge.[2] In particular, it seems that Ted's inferred belief is not luckily true like

Chandler's inferred belief. Indeed, this looks like a very secure way of forming a true belief. But if one can't in general gain knowledge from inferring a true belief from a false (even if justified) belief, then why can one gain such knowledge in this scenario?

Or consider this scenario:

> *President*
> CNN breaks in with a live report. The headline is 'The President is speaking now to supporters in Utah'. Fritz reasons that since the President is in Utah, he is not attending this morning's NATO talks in Brussels. The President is not in Utah, however, but Nevada—he is speaking at a 'border rally' at the border of those two states and the speaking platform on which he is standing is in Nevada (it is the crowd listening to the speech that is in Utah).[3]

As with Handout, the subject is forming a true belief by responsible inferring it from what is, unbeknownst to them, a false belief. Nonetheless, the resulting belief doesn't seem to be luckily true in the way that the corresponding belief in Detective is. Relatedly, it also seems to amount to knowledge.

It thus seems that there is an important epistemic difference between cases like, on the one hand, Handout and President, and cases like, on the other hand, Detective, despite their structural similarities. This therefore calls for a diagnosis of what grounds this epistemic difference.

2. Call cases like Handout and President, where knowledge seems to result from an inference from a falsehood, *knowledge from falsehood* cases. Some commentators have disputed whether such cases really are instances of deriving knowledge from falsehood, at least to the extent that a falsehood is playing a necessary role in the resulting inferential knowledge anyway.[4] For example, in a case like Handout, one could argue that what's really load-bearing in the relevant inference is not so much the subject's (false) belief that there are 53 people in the room as their more general (and true) conviction that there are a lot less than 100 people in the room. In contrast, while I would grant that some cases might be amenable to such an interpretation, I'm not convinced that this strategy has general application. Accordingly, I propose to proceed by granting that there are genuine cases of knowledge from falsehood. The challenge I face is thus to account for why knowledge is *bona fide* in such cases when it is not generally possible to gain knowledge by inferring it from a falsehood.

One straightforward way of explaining why such cases are possible is to appeal to the safety condition on knowledge.[5] Many epistemologists hold that a necessary condition for knowledge is that the target belief is formed on a safe basis, such that it could not have very easily been a false belief, so formed.[6] With this constraint on knowledge in mind, one has a plausible way of

distinguishing between the usual cases involving inference from falsehood that don't result in knowledge (like Detective), and the more specific class of cases like Handout which, it is claimed, do result in knowledge. In particular, while in general one cannot gain a safe belief from inferring it from a false belief, there are some cases, like Handout, where this is possible, and which are thus compatible with the acquisition of knowledge.

That reasoning from a false belief is generally an unsafe basis for belief ought to be uncontentious. The Detective case illustrates why. While Chandler formed a true belief via the target inference, and did so by inferring it from a justified belief, the inferred belief is not safe. Via this reasoning Chandler could have very easily formed a false belief, given that it is just a coincidence that Jones' blood type happens to be the same as the murder's. Or consider the Gettier case described above. Although the inferred belief is true, it could very easily have been false, given how it was formed, such as the scenario in which Jones hadn't recently travelled to Barcelona.

Notice too that it doesn't matter in this regard whether the inferred belief happens to be a necessary truth, as it would still be the case that the subject's inferred belief would be unsafe. For example, imagine that in the Gettier case the second disjunct happened to be necessarily true, such that the disjunction as a whole was also necessarily true. This is still compatible with a belief in this disjunction being unsafe since, as safety is usually understood, what's important is that one's basis for belief couldn't have easily resulted in a false belief, where the proposition believed needn't be the same proposition that the subject actually believes. For instance, suppose one forms one's belief that $2 + 2 = 4$ by flipping a coin. There is obviously no close possible world where one forms that same belief on the same basis and believes falsely, since there is no possible world where the target proposition is false. But there is a close possible world where that basis for belief results in a false belief, such as the possible world where flipping a coin leads one to believe that $2 + 2 = 5$. In this way, safety theorists are able to explain why even beliefs in necessary truths can be unsafe.[7] Accordingly, even if our hero had inferred a disjunction with a random claim as the second disjunct that happened to be (unbeknownst to our agent) a necessary truth, his belief would still be unsafe, as there will be a close possible world where that way of forming a belief in a disjunction results in a false belief.

With the foregoing in mind, we are now in a position to draw a contrast between typical instances of inferences from falsehood, like Detective, and knowledge from falsehood cases like Handout and President. In Handout, while the initial belief might be slightly inaccurate, the belief inferred on this basis is nonetheless safe. In particular, Ted could not very easily form a false belief

via the relevant inference, given that the number of handouts that he has is comfortably sufficient for the actual number of people in the room. The crux of the matter is that false beliefs that are approximately true can provide safe bases for true beliefs which aren't affected by the relevant margin of error.

The same is also true of the inferred belief in the President case. Although the initial belief is false, it is nonetheless sufficiently in the ballpark of the truth to ensure that the inferred belief is safe. Whereas the Handout case ensures this by having a numerical distance between the number at issue in the original belief and the threshold at issue in the inferred belief, the President case ensures this by there being a geographical distance. Given how far Brussels is from the region where the President is located, it follows that there is no close possible world where he makes it across to Brussels for that morning's NATO meeting, which is why Fritz's inferred belief is safe regardless. The geographical distance in this case, like the numerical distance in the Handout scenario, ensures there is the required modal distance to underwrite the safety of the inferred belief.

This way of explaining knowledge from error cases can also account for why analogous scenarios where there isn't the same modal distance wouldn't generate knowledge, because the resulting belief would not be safe. Consider this variant on the Handout scenario:

> *Handout**
> Ted* needs to determine how many handouts he will need for his talk. Carefully counting the number of people present he forms the belief that there are 99 people in the room. He accordingly infers that the 100 handout copies that he has will be sufficient. But Ted* miscounted, and there are in fact 98 people in attendance.

I take it that there is no temptation to ascribe knowledge to Ted* in this scenario. The only difference, however, is the numerical gap between the counted number of people in the room and the threshold number at issue in the inferred belief. The natural explanation of why this change makes such a difference to the case is that with this numerical gap reduced to a sliver it is now very easy for Ted* to end up with a false belief by making this inference. In particular, it could have easily been the case that Ted*'s counting of the people in the room was off such that his inferred belief about his 100 handout copies being sufficient is false. Ted* thus has an unsafe basis for belief, in contrast to Ted. We can easily imagine a parallel variant of President where similar points apply (e.g., where the subject infers that the President is not attending this morning's NATO meeting in Oregon), and thus where the inferred belief is unsafe.

3. Since safety is only held to be necessary for knowledge, it obviously doesn't follow that the inferred belief in the Handout case, or cases like it, amount to knowledge, but at least we are able to point to a condition on knowledge which is present here but absent in the relevant Gettier-style cases.

Moreover, insofar as the rationale behind safety is to accommodate intuitions about the incompatibility of knowledge with veritic epistemic luck or high levels of epistemic risk, then we can unpack the claim being made here in these terms.[8] Consider *anti-risk epistemology*, for example. This appeals to a modal account of risk which holds that, roughly, we should evaluate risk by considering the modal closeness of the target risk event. So, for example, if the target risk event associated with plane travel is dying in a plane crash, then to treat such a mode of transport as high risk is to contend that there is a close possible world where someone who takes a flight ends up dying in a plane crash.[9] With the modal account of risk applied to knowledge, and the target epistemic risk event identified as the formation of a false belief, we can see how anti-risk epistemology provides a rationale for the safety condition on knowledge. Knowledge excludes high levels of epistemic risk in that it is incompatible with one forming one's true belief on a basis that could very easily have led to the epistemic risk event of forming a forming a false belief. *Ergo*, knowledge entails safety.[10]

With this rationale in mind, we have a way of explaining why knowledge might be in general incompatible with inference from falsehood but not universally so. In short, the explanation is that while inferring a true belief from a falsehood would generally make one's belief subject to high levels of epistemic risk—in that the epistemic risk event of forming a false belief is modally close—this is not always the case. In particular, there is a particular class of cases, exemplified by Handout and President, such that a true belief formed on this basis is not subject to a high level of epistemic risk, in that the epistemic risk event is not modally close.

Relatedly, we can use anti-risk epistemology explain why there would be norms of epistemic responsibility what would be applicable to these inferences in the way articulated above. As epistemically responsible believers, Ted and Fritz, the subjects in the Handout and President cases, are surely aware of how the inference they are undertaking is by its nature epistemically low risk, and this informs their willingness to make this inference. So long as their initial belief is roughly correct, their inferred belief will be bound to be safe. In contrast, given the moderate nature of their epistemic support for the entailing belief, as epistemically responsible believers they would be wary about undertaking corresponding inferences that lack the modal distance to the target risk event, and which are thus epistemically high risk. Counting large numbers of people, even carefully, is prone to

error, and as a responsible believer Ted will be aware of this. Relatedly, catching a single news item on the TV is a relatively fallible way of determining the current location of the President. In both cases this level of fallibility is compatible with an epistemically responsible subject forming a belief in the target proposition, but not with the agent drawing inferences from it where there is no margin for error, and which are thus epistemically high risk. This is why Ted*'s willingness to infer that the number of handouts is sufficient indicates that he is not an epistemically responsible believer, as he ought to be sensitive to the fact that this is an epistemically high risk inference. Part of what it is to be epistemically responsible in one's inferences is thus to be alert to the extent to which the inference is epistemically risky.

4. We thus have a potential rationale, in terms of the safety condition—and thus the anti-luck/anti-risk epistemology that motivates the safety condition—for distinguishing between genuine cases of knowledge from error from corresponding inferences from falsehood where the inference does not generate knowledge. So are we then home and dry in terms of accounting for the phenomenon of knowledge from error? I don't think so.

In order to see why, we need to note that this account of knowledge from falsehood as it stands implies that any case where a subject infers a true belief from a falsehood where the inferred belief is safe ought to be in the market for knowledge. After all, the reason why knowledge was lacking in a case like Detective is that the inferred true belief is unsafe, unlike the inferred true belief in a case like Handout. Accordingly, provided that the true belief is safe, and there is nothing independently epistemically amiss with the target belief, then why wouldn't it amount to knowledge? Think of this in terms of epistemic risk. If the reason why knowledge is lacking in a case like Detective is that the level of epistemic risk is too high, but safety ensures that levels of epistemic risk are low, then why wouldn't a safe true belief formed on the basis of an inference from falsehood (and which is not epistemically amiss in any other way) not be an instance of knowledge, just like the Handout case?

The reason why I am labouring this point is that there are instances with a similar structure to knowledge from error cases where the inferred true belief is safe but where it does not amount to knowledge. Consider the following scenario:

> *Temperature*
> Thom is tasked with doing regular temperature readings of the industrial oven in operation at his workplace. For this he uses a digital thermometer that provides exact readings in Centigrade. He then

> converts the readings to Fahrenheit and writes down on a chart which 10-degree temperature range that reading belongs to (e.g., 220°F-229°F, etc.,). Unbeknownst to Thom, the thermometer is broken, and the readings it is generating are completely random and hence almost always false. (Thom has no idea what readings he should be getting, and so he remains oblivious to the random nature of the readings). Also unbeknownst to Thom, however, there is someone secretly observing him for whom it is crucially important that the results that Thom enters into the chart are correct. As a result, whenever this person sees Thom take a reading from the thermometer he adjusts the temperature in the oven to ensure that it is comfortably within the appropriate Fahrenheit range by the time that Thom enters the converted reading into his chart.

Thom is thus responsibly drawing inferences from false beliefs (regarding the temperature readings in Centigrade from the faulty digital thermometer) that result in true beliefs about which Fahrenheit temperature band corresponds to the temperature in the industrial oven. Moreover, these inferred beliefs are entirely safe, given how they are formed, as Thom could not easily form a false belief in this manner; indeed, he is to all intents and purposes *guaranteed* to form a true belief in this regard because of the intervention of the hidden agent.

This case is thus unlike a standard case involving an inference from a false belief, in that the inferred belief is safe. Crucially, however, this scenario is also very different from the kind of knowledge from error cases that we looked at above, Handout and President. In particular, Thom's inferred belief doesn't seem to be in the market for knowledge at all, even despite its safety. This is because the safety of Thom's belief has nothing whatsoever to do with his exercise of cognitive agency but is rather entirely due to the intervention of the hidden agent.

Compare Thermometer in this regard with a case like Handout where the subject's inferred safe belief does amount to knowledge. Although Ted is inferring a true belief from a false one, the safety of the inferred belief is significantly due to his exercise of cognitive agency—i.e., the cognitive agency involved in his false-but-approximately-true original belief and his inference to the entailed true belief, the truth of which lies a comfortable numerical distance from the original belief. Indeed, as we noted above, the intuition that Ted is a knower in this case depends on treating him as the kind of subject who wouldn't have made this inference had it concerned a threshold that was close to the numerical value of his entailing belief, and which would thus have resulted in an unsafe (albeit true) belief. The crux of the matter is that Ted's cognitive agency is playing an important explanatory role with regard to his safe cognitive success. This is not replicated in the Temperature case, however, as here the safety of the inferred belief is entirely attributable to the intervention of the helper rather than the cognitive agency of our hero.

5. We thus have three types of scenario that concern us. First, there are standard cases involving inferring a true belief from a falsehood, such as inferential Gettier cases or Detective, where the subject responsibly infers an unsafe true belief from (unbeknownst to them) a false belief, and hence lacks knowledge. Second, there are genuine knowledge from error cases, like Handout and President, where the subject responsibly infers a safe true belief from (unbeknownst to them) a false belief, but where the safety of the true belief is significantly creditable to the subject's cognitive agency. The subject thereby gains knowledge, even despite the false belief involved. Finally, we have a third kind of case, illustrated by Temperature, where the subject responsibly infers a safe true belief from (unbeknownst to them) a false belief, but where the safety of the true belief is not significantly creditable to the subject's cognitive agency. I'm suggesting that in these third kind of cases the subject does not gain knowledge. In particular, if the foregoing is correct, then it is a mistake to think that what differentiates genuine knowledge from error cases from other scenarios where an inference from a false belief doesn't lead to knowledge is simply whether the inferred belief is safe, as there is a further factor that we need to be alert to here (but which is easily overlooked). This is whether the safety of the subject's belief is significantly attributable to her manifestation of cognitive agency.

Relatedly, insofar as the rationale for safety is an anti-risk (or anti-luck) epistemology, one can't differentiate the second and third category of cases by appealing to this rationale alone, as they both concern a subject forming a safe true belief in the inferred proposition. Nonetheless, this way of explaining the difference is on the right lines. What it is missing is the further explanatory relation that the subject's safe cognitive success should satisfy if it is to amount to knowledge. This point is obscured in genuine knowledge from error cases as the subject's cognitive success is both safe *and* the safety of this cognitive success is significantly attributable to her manifestation of cognitive agency. As we have seen, however, where the latter is absent there is no longer any temptation to ascribe knowledge, which indicates that it is this aspect of knowledge that is crucial to understanding why knowledge is present in such cases.

Moreover, we can also adapt the motivation offered by anti-risk epistemology to explain why there is knowledge in genuine knowledge from error cases. It is not enough for knowledge that one's cognitive success is not subject to high levels of epistemic risk (i.e., safe), as it is also required that this feature of one's cognitive success is significantly attributable to one's cognitive agency. Elsewhere I have argued for this general way of thinking about knowledge under the description *anti-risk virtue epistemology*, in that it maintains that knowledge involves an interplay between the

exclusion of epistemic risk (as represented by the necessity of the safety condition) and the manifestation of cognitive ability (as represented by the required explanatory relation between one's safe cognitive success and one's cognitive agency).[11] In particular, what is important is that the immunity of one's true belief to high levels of epistemic risk bears an appropriate explanatory connection to one's manifestation of cognitive ability. This is why knowledge is lacking in cases like Temperature, even though the subject's responsibly inferred true belief is safe. But it is also why knowledge is present in genuine cases of knowledge from error, in that while the inference is made from a false belief, the resulting belief is not only safe, but safe in a manner such that it is significantly attributable to the agent's cognitive agency.

6. We have claimed that knowledge from error is a genuine phenomenon, and also explained how it can be possible, given that in general one does not acquire knowledge by inferring a true belief from a falsehood. Our explanation primarily concerned an appeal to the safety of the beliefs so formed (and the general unsafety of beliefs inferred from falsehoods), and the underlying anti-risk motivation for safety. As we saw, however, this story needs to be complicated, for what is in fact doing the work of distinguishing knowledge from error cases from parallel cases where knowledge is not acquired is not just the safety of the inferred belief, but also the fact that this safe true belief is significantly attributable to the subject's manifestation of cognitive ability. We thus get an explanation for why there can be knowledge from error via appeal to a distinctive and independently motivated way of thinking about knowledge known as anti-risk virtue epistemology.[12]

REFERENCES

Ball, B., & Blome-Tillmann, M. (2014). 'Counter Closure and Knowledge Despite Falsehood', *Philosophical Quarterly* 64, 552-68.

Bernecker, S. (2020). 'Against Global Method Safety', *Synthese* 197, 5101-16.

Black, T. (2011). 'Modal and Anti-Luck Epistemology', *Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, 187-98, London: Routledge.

Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* 23, 121-23.

Klein, P. (2008). 'Useful False Beliefs', *Epistemology: New Essays*, (ed.) Q. Smith, 25-61, Oxford: Oxford University Press.

Fitelson, B. (2010). 'Strengthening the Case for Knowledge from Falsehood', *Analysis* 70, 666-69.

Luzzi, F. (2014). 'What Does Knowledge-Yielding Deduction Require of its Premises?', *Episteme* 11, 261-75.

Montminy, M. (2014). 'Knowledge Despite Falsehood', *Canadian Journal of Philosophy* 44, 463-75.

Pritchard, D. H. (2002). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.

——— (2004). 'Epistemic Luck', *Journal of Philosophical Research* 29, 193-222.

——— (2005). *Epistemic Luck*, Oxford: Oxford University Press.

——— (2007). 'Anti-Luck Epistemology', *Synthese* 158, 277-97.

——— (2012*a*). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109, 247-79.

——— (2012*b*). 'In Defence of Modest Anti-Luck Epistemology', *The Sensitivity Principle in Epistemology*, (eds.) T. Black & K. Becker, 173-92, Cambridge: Cambridge University Press.

——— (2014). 'The Modal Account of Luck', *Metaphilosophy* 45, 594-619.

——— (2015*a*). 'Anti-Luck Epistemology and the Gettier Problem', *Philosophical Studies* 172, 93-111.

——— (2015*b*). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*, Princeton, NJ: Princeton University Press.

——— (2015*c*). 'Risk', *Metaphilosophy* 46, 436-61.

——— (2016). 'Epistemic Risk', *Journal of Philosophy* 113, 550-71.

——— (2017*a*). 'Anti-Risk Epistemology and Negative Epistemic Dependence', *Synthese* (Online First, DOI: https://doi.org/10.1007/s11229-017-1586-6).

——— (2017*b*). 'Knowledge, Luck and Virtue: Resolving the Gettier Problem', *The Gettier Problem*, (eds.) C. Almeida, P. Klein & R. Borges, 57-73, Oxford: Oxford University Press.

——— (2020). 'Anti-Risk Virtue Epistemology', *Virtue Epistemology*, (ed.) J. Greco & C. Kelp, 203-24, Cambridge: Cambridge University Press.

Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.

Sainsbury, R. M. (1997). 'Easy Possibilities', *Philosophy and Phenomenological Research* 57, 907-19.

Schnee, I. (2015). 'There is no Knowledge from Falsehood', *Episteme* 12, 53-74.

Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.

Warfield, F. (2005). 'Knowledge from Falsehood', *Philosophical Perspectives* 19, 405-16.

Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.

NOTES

[1] This example is originally due to Warfield (2005, 407-08).
[2] See, for example, Warfield (2005), Klein (2008), Fitelson (2010), and Luzzi (2014).
[3] This case is also originally due to Warfield (2005, 408).
[4] For some discussions of this way of responding to putative knowledge from error cases, see Warfield (2005), Ball & Blome-Tillmann (2014), Montminy (2014), and Schnee (2015).
[5] The first to notice this point was, I believe, Warfield (2005, 414-15), albeit not quite in these terms (i.e., while he doesn't mention safety specifically, this does seem to be the kind of condition on knowledge that he has in mind).
[6] We will expand on the notion of safety in a moment. For some of the main defences of safety, see Sainsbury (1997), Sosa (1999), and Williamson (2000). I develop my own account of safety in Pritchard (2002; 2005, *passim*; 2007; 2012*a*; 2012*b*; 2015*a*).
[7] For further discussion of this point, see Pritchard (2007*a*; 2012*a*; 2012*b*).
[8] In earlier work I offered a detailed defence of the idea that the safety condition on knowledge should be understood as excluding veritic epistemic luck—see Pritchard (2004; 2005, *passim*; 2007; 2012*a*; 2012*b*). In more recent work, I have further developed this claim by focusing on epistemic risk specifically—see Pritchard (2015*c*; 2016; 2017*a*; 2020).
[9] For further defence of the modal account of risk, see Pritchard (2015). See also the closely related modal account of luck, as articulated, for example, in Pritchard (2014).
[10] Indeed, anti-risk epistemology—and, for that matter, anti-luck epistemology—motivates a particular way of interpreting the safety principle, including *inter alia* the idea that safety is not proposition-specific, and so can accommodate unsafe beliefs in necessary truths, as noted above. See Pritchard (2007; 2012*a*; 2012*b*; 2016; 2020).
[11] See especially Pritchard (2020). Note that anti-risk virtue epistemology is a refinement of my earlier account of knowledge, *anti-luck virtue epistemology*, and as such shares many of its core structural features. For further discussion of the latter, see Pritchard, Millar & Haddock (2010, ch. 4) and Pritchard (2012*a*).
[12] For helpful discussion of topics related to this paper, I am grateful to Sven Bernecker and Bin Zhao.