THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Anti-risk virtue epistemology

Citation for published version:

Pritchard, D 2020, Anti-risk virtue epistemology. in C Kelp & J Greco (eds), *Virtue Theoretic Epistemology: New Methods and Approaches*. Cambridge University Press, pp. 203-224.
<https://doi.org/10.1017/9781108666404.010>

Digital Object Identifier (DOI):

[10.1017/9781108666404.010](https://doi.org/10.1017/9781108666404.010)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Virtue Theoretic Epistemology

Publisher Rights Statement:

This material has been published in *Virtue Theoretic Epistemology: New Methods and Approaches* edited by Christoph Kelp and John Greco. This version is free to view and download for personal use only. Not for re-distribution, re-sale or use in derivative works. © Duncan Pritchard.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

*Anti-risk Virtue Epistemology**Duncan Pritchard***9.1 Introductory Remarks**

For more than a decade now I have been defending a particular kind of methodology as regards the theory of knowledge, which I call *anti-luck epistemology*.¹ In essence, this argues that we should take seriously the platitude that knowledge excludes luck. This means both giving an account of the nature of luck and also articulating the precise manner in which knowledge is incompatible with it. Once one puts these two components together, the thought runs, then one will have formulated, on a principled basis, the anti-luck condition on knowledge.²

Such a methodology has many advantages, not least because it offers a more systematic way of approaching the problems posed by the theory of knowledge. For one thing, we are now not simply formulating a condition and then testing it against various counterexamples to see how it fares relative to competing conditions, in a piecemeal manner. For another, the condition we come up with will cover all cases where it matters that knowledge and luck are incompatible, as opposed to the more standard approach of trying to find an anti-Gettier condition and then seeing whether, hopefully, it will also have application to other kinds of cases that trade on luck (such as lottery cases, for example).

In previous work I've made a number of claims about anti-luck epistemology. For instance, I've argued that it not only motivates the safety condition on knowledge over other competing modal conditions on knowledge (such as sensitivity), but also gives us an explanation of how

¹ See, for example, Pritchard 2004, 2005, 2007, 2012a, 2012b, 2015a, and Pritchard, Millar and Haddock 2010: chs. 1–4.

² Another aspect of the anti-luck epistemology methodology that I emphasise in my presentations of it is its appeal to the empirical literature on luck (and, relatedly, risk) ascriptions. See, for example, Pritchard 2007. See n. 9 for a list of some of the relevant empirical literature. For a survey of some of this literature, see Pritchard and Smith 2004.

this condition should be interpreted – an interpretation that avoids some familiar problems that have been levelled at it.³ I've also argued that anti-luck epistemology not only provides us with an account of the anti-luck condition on knowledge, but that we can use this as a basis to offer a complete virtue-theoretic account of knowledge, one that does justice to both the platitude that knowledge excludes luck and the related platitude – the *ability intuition*, as I call it – that knowledge entails the manifestation of relevant cognitive ability. I describe the resulting theory of knowledge as *anti-luck virtue epistemology*.⁴

My concern here, however, is not with further defending and articulating anti-luck epistemology (or anti-luck virtue epistemology). Rather, I want to explain a recent shift in my position. For while I think that anti-luck epistemology and anti-luck virtue epistemology are broadly on the right lines, I have become convinced that both need to be refined in an important way. In particular, I now argue that the relevant notion that we should focus on as epistemologists (and not just as theorists of knowledge, as I explain below) is in fact *epistemic risk* rather than epistemic luck.⁵ Anti-luck epistemology is thus to be replaced with *anti-risk epistemology*. Moreover, and crucially for our purposes, we will see that this also means offering a new virtue-theoretic account of knowledge: *anti-risk virtue epistemology*.

9.2 Luck and Risk

In order to understand the transition from anti-luck epistemology to anti-risk epistemology, we first need to appreciate some of the subtle – though as we will see, important nonetheless – differences between luck and risk. Luck and risk are very closely related notions. Indeed, it is quite common for psychologists to study risk and luck ascriptions together.⁶ It is easy to

³ See especially Pritchard 2007, 2015a. For example, I've claimed that anti-luck epistemology motivates a version of the safety principle that evades the dilemma posed for safety by Greco (2007). I've also claimed – e.g. in the works just cited – that it is better placed to handle a problem posed to safety by necessary, or at least modally stable, truths (although as we will see below, I think anti-risk epistemology is in fact on stronger ground in this regard). Versions of the safety principle had previously been defended by such figures as Sainsbury (1997), Sosa (1999), and Williamson (2000). For a comparison of safety and the competing sensitivity principle in the context of an anti-luck epistemology, see Pritchard (2008).

⁴ See especially Pritchard 2012a and Pritchard, Millar, and Haddock 2010: chs 1–4). See also Pritchard 2017b.

⁵ See especially Pritchard 2016b. See also Pritchard 2017a.

⁶ See, for example, Teigen 1998a. For discussion of this point, and further elaboration on the empirical literature in this regard, see Pritchard and Smith 2004.

see why. Suppose you are a journalist working in a war zone and a sniper's bullet whizzes past your ear, narrowly missing you. You were lucky not to be shot. There was also a high risk of you being shot. Or suppose you are on a plane that crash lands, albeit safely, into the Hudson River. You were lucky not to die. You were also at a high risk of dying. And so on. Cases like this show that, often at least, luck and risk go hand-in-hand. If it is lucky that such-and-such event failed to happen, then that tends to correlate with it being the case that there was a high risk of it happening.

The close relationship between the two notions means that it is easy to overlook their differences. And yet, as we will see, the differences are in fact important, and important in a way that is particularly relevant to epistemology. I've argued at length elsewhere that we need a *modal account of luck*. This holds, roughly, that a lucky event is an event that actually occurs but which doesn't occur in close possible worlds where the relevant initial conditions for that event remain the same.⁷ Narrowly avoiding being shot by a sniper's bullet is lucky because there is a close possible world where, keeping the relevant initial conditions fixed (e.g. someone is still trying to shoot you, and so forth), you do get shot. In contrast, there is nothing lucky about not being shot in normal circumstances where there are no snipers around, since there are no close possible where one gets shot. Notice too how the modal account of luck can also account for how luck comes in degrees. *Ceteris paribus*, narrowly avoiding a sniper's bullet by inches is luckier than avoiding it by several feet, since the non-obtaining of the target event is modally closer in the first scenario.⁸

Or consider, for example, a paradigm case of a lucky event: a lottery win. On the modal account this is a lucky event because there are very close possible worlds where the relevant initial conditions for that event remain the same (e.g. you continue to buy a lottery ticket, the lottery remains free and fair, and with the same long odds of winning, and so on), but you fail to win. The lottery win is an interesting case because it also illustrates why one wouldn't want a straightforwardly probabilistic account of luck. This might be initially surprising, since a probabilistic account of luck, on the face of things, looks promising. Aren't lucky events those events that are probabilistically unlikely to happen? So, for example, isn't the lottery win a lucky event because it was a low probability event? This is

⁷ See especially Pritchard 2005: ch. 5, 2014. Note that luck doesn't just apply to events, but I will henceforth focus on events in order to simplify my discussion (and I will be making the same simplifying assumption in my discussion of risk).

⁸ The modal account of risk that I describe below inherits this feature.

too quick, however, for consider the person who plays the lottery but loses. It is also true of them that they could very easily have won – all it would have taken, after all, is that a few coloured balls fall in a different configuration. And yet the odds massively support the possibility that they would lose. The crux of the matter is that there can be very low-probability events (in this case winning the lottery) that can nonetheless be modally close. But in such cases we still treat the event as being down to luck (bad luck, in the case of losing the lottery), something that is widely confirmed by the empirical literature on luck ascriptions.⁹ Indeed, isn't that just why we play the lottery (but don't place bets on scenarios with similar long odds but which are modally far-fetched, like the North Korean soccer team winning the World Cup)?¹⁰

The modal account of luck, while obviously not without its detractors, is certainly the leading view in the literature.¹¹ Given the close connections between luck and risk, a modal account of risk would be similarly appealing, such that, roughly, our concern in making a risk assessment is the modal closeness of the target event (rather than, for example, its probability).¹² But once one develops such an account, however, one starts to notice the small, but important, differences between the notions of luck and risk.

The first big difference is that luck is about the non-obtaining of the target event. To have a lucky lottery win entails that you could have very easily lost the lottery drawing, in relevantly similar circumstances. To have luckily avoided being shot by a sniper while working as a journalist in a war zone entails that you could have easily been shot by that sniper. And so on. The important point is that our focus is on the closeness of the possible

⁹ See, for example, Teigen 1995, 1996, 1997, 1998a, 1998b, 2003; Tetlock 1998; Kahneman and Varey 1990; and Tetlock and Lebow 2001. For a survey of some of this empirical literature, see Pritchard and Smith 2004.

¹⁰ The UK national lottery used to trade on this feature in its advertising campaign. They had the slogan 'It could be you!' and featured someone being picked out as a winner by a God-like finger. Clearly this is not the 'could' of probability, as in this sense it *couldn't* be you, as the odds are so astronomically long. Rather it is the 'could' of modal closeness: if you play the lottery, then someone, just like you, will be a winner; all that is required is that a few numbered balls fall in a slightly different configuration. (Note too that the UK has very strict advertising laws, and it is unlikely they would have got away with such a campaign slogan had the probabilistic reading been the only one available).

¹¹ To be fair, this is largely because for a while it was the *only* full-fledged philosophical account of luck in the literature. For some detractors/proponents of rival (or at least distinct) views, see Riggs 2007, 2009; Lackey 2008; Levy 2011; Broncano-Berrocal 2015; and Coffman 2015.

¹² Interestingly, the dominant account of risk in the literature is probabilistic – see, for example, Hansson 2014. See also Hansson 2004. I argue against the probabilistic conception of risk in Pritchard 2015b, where I also articulate the modal account of luck.

world where the target event that actually occurred failed to occur (while keeping the relevant initial conditions fixed).

Risk is slightly different in this regard. When we make assessments of risk, we have in mind some specific *risk event* that we want to avoid. So, for instance, when flying we might make an assessment of risk relative to the risk event of the plane crashing. Or, when mountain climbing, we might make an assessment of risk relative to the risk event of falling off the mountain. On the modal account of risk, these assessments will be concerned with the modal closeness of the risk event (keeping all relevant initial conditions fixed), such that the closer it is, the riskier it is.¹³ But won't that make the modal account of risk exactly like the modal account of luck? Interestingly, no.

The reason for this is that we can vary our risk assessments by varying the risk event that we focus upon. So, for example, we may judge that in flying there is a far greater risk of the plane crashing at take-off and landing than there is while in flight. That is, we might judge that the former scenario is modally closer, and hence riskier, than the latter scenario (which it is, I gather). Changing our target risk event thus changes the level of risk involved in the target event (in this case, taking the plane). There is no limit to the range of risk events that we can focus upon, and indeed we can even imagine complex risk assessments that involve multiple risk events (though to keep matters simple I will just focus on assessments involving a single risk event).

Notice the difference with luck. With luck, we first identify a target event, and then we consider the modal closeness of that event not obtaining. Once the event has been determined, then there is no further parameter that we need to consider in evaluating whether it is lucky. With risk, however, there *is* a further parameter, since we also need to select the risk event that we wish to focus upon. I think this point gets overlooked because of course we can change the direction of focus with luck too, by considering a different target event. Rather than focusing on the event of avoiding being killed by a sniper's bullet, for example, we could focus on the event of avoiding being shot (but not necessarily killed). This will undoubtedly make a difference to how lucky we think the target event is. But the point is that with luck, but not risk, once you've selected your target event, there is no further parameter that needs to be fixed to determine whether the event is lucky (as there is with risk). This detail might seem so small as to be insignificant, but as we will see it is in fact very important to resolving a (hitherto unnoticed) difficulty facing anti-luck

¹³ See Pritchard 2015b for further discussion of the modal account of risk.

epistemology. It also enables an anti-risk epistemology to offer a broader set of epistemic assessments than anti-luck epistemology.

There are further differences between luck and risk that we should flag. One interesting difference is that risk assessments are essentially forward looking, while luck assessments are essentially backward looking. When we judge that we were lucky to have avoided the sniper's bullet, we are making this assessment from the perspective where things turned out well, luckily for us (as in 'phew, I was lucky there'). In contrast, when we judge that we were at a high risk of being shot we are taking the forwards-looking perspective of considering our situation *prior* to the sniper firing, and responding to the clear danger that we are facing.¹⁴ As we will see, this has important diagnostic implications for our understanding of post-Gettier epistemology.

A further point in this regard is that our interest in luck – *bad* luck, anyway (which happens in any case to be the main concern of epistemologists) – is usually motivated by a concern to eliminate risk, rather than vice versa. Why is it significant that one was lucky not to fall off the mountain? Well, ordinarily at least, because one was concerned that one was at such a high risk of falling (and dying). Why is it significant that one was lucky not to die in the plane crash? Well, ordinarily at least, because one was concerned at the high level of risk of dying in that plane crash. The point is that while luck and risk tend to go hand-in-hand, our deeper concern is with the latter rather than the former, such that our interest in (bad) luck typically signals a concern to eliminate risk. As we will see, this relationship between luck and risk also has important diagnostic implications that an anti-risk epistemology can exploit to its advantage.

There are other differences between luck and risk, but many of them aren't salient for our current purposes. I will just mention one, since it's interesting in its own right, which is that luck is clearly double valenced, in that it comes in both good and bad forms (though, as just noted, in epistemology it is usually just the bad epistemic luck that we are interested in). Risk, in contrast, tends to concern negative outcomes, in that the risk event is something that we want to avoid (such as a plane crash). Even when risk adds value to an activity, as with extreme sports or aesthetic risk, the risk event is still something negative (serious injury, ruining the performance).¹⁵

¹⁴ This difference came to light in discussion with Jesus Navarro.

¹⁵ I discuss 'positive' kinds of risk of this kind, with a particular focus on aesthetic risk, in Pritchard 2018. I offer a more detailed description of the modal account of risk in Pritchard 2015b. See also Pritchard 2017c, where I apply the modal account of risk to epistemological issues in the philosophy of law.

9.3 Anti-risk Epistemology

So, how will an anti-risk epistemology differ from an anti-luck epistemology? We first need to understand something that unsettled me about anti-luck epistemology from the off (though, oddly, no-one else clocked this theoretical lacuna, so far as I know). Recall that the methodology of anti-luck epistemology involves putting together a theory of luck with an account of the specific way in which knowledge excludes luck. We've encountered the modal account of luck already. Sparing the reader the details, the specific way in which knowledge is incompatible with luck is *veritic epistemic luck*, whereby it is a matter of luck that one's belief is true, given how it is formed.¹⁶ Given the modal account of luck, this means that we have a condition on knowledge such that the non-obtaining of the target event (true belief) could very easily have occurred. But, strictly speaking, this entails that our focus shouldn't be on the closest possible world where one believes *falsely* on the same basis, but rather the closest possible world where one *fails to form a true belief* on the same basis, which is a disjunction of possible worlds where one forms a false belief and where one fails to form a belief at all.

It is the latter part of the disjunction that is troubling, as it doesn't seem to be necessarily knowledge-undermining to form a true belief on a basis whereby one could very easily have not formed a belief at all on that same basis. Perhaps one is just a cautious believer? Indeed, isn't being a cautious believer often a sign that one forms one's beliefs in an epistemically conscientious way, and hence that one is more likely to be a knower? But how is one to restrict the view, in a principled fashion, according to anti-luck epistemology? My thought, way back in Pritchard (2005), was that the fact that one had initial conditions built into the modal account of luck, and hence had an independent rationale for a basis-relative account of the anti-luck condition (safety), would offer a response to this problem. The idea was that so long as one takes basis-relativity seriously, then this problem disappears because in the relevant cases one's basis is never such as to allow mere non-belief.

Consider, for example, a case in which one believes that one's lottery ticket is a loser because one reads the result in a national, well-established, newspaper. We now need to examine the closest possible worlds where one

¹⁶ Veritic epistemic luck is thus, as I put it, a *malign* form of epistemic luck, in that it is incompatible with knowledge, as opposed to *benign* forms of epistemic luck, like the *evidential epistemic luck* that one is lucky to have the evidence that one does. For the details, see Pritchard 2005, *passim*.

has the same basis for belief as in the actual world. This is going to entail reading the lottery result in a national, well-established, newspaper. But if that's right, then there won't be a close possible world where the basis is the same but one fails to form any belief at all, as reading the result in the newspaper is going to lead one to form *some* judgement about whether one's ticket is a winner or a loser.

The problem, however, is that while it's true that appealing to the basis-relativity of safety will deal with most cases, it does not deal with all of them. Consider, for instance, a subject who has a sound memorial basis for their belief that the Battle of Hastings was in 1066. Suppose, however, that although our subject is accordingly confident of this belief in the actual world, there are close possible worlds where they is disposed to doubt themselves. It is just a psychological fact about this subject that their confidence in what they believe is quite variable, even though there is no epistemic basis for this variability (it is not as if, for example, they have reasons to doubt their memory).¹⁷ It will now be true of such a subject that in close possible worlds they will have the same basis for belief as in the actual world and yet will not believe the target proposition on this basis. It would thus follow, on the interpretation of safety motivated via anti-luck epistemology, that our subject lacks knowledge. But that doesn't seem like the right result at all. After all, in all close possible worlds where they form a belief on this sound memorial basis their belief will be true. So why should their psychological reticence to form beliefs on this basis in some close possible worlds deprive them of knowledge? Indeed, as noted above, one could argue that, *ceteris paribus* anyway, cautious believers like our subject are to be preferred, from an epistemic point of view, to those who are far too quick to form beliefs.

Could we get around this problem by stipulating that a basis for belief should be understood in such a way that it entails that the subject forms a belief? That would certainly fix the problem, but it is hard to see what motivation there would be for this requirement, other than the post hoc one of evading this difficulty. On any plausible rendering of the basing relation, one can surely be in possession of the basis and, in principle anyway, form *any* doxastic attitude about the target proposition.¹⁸

¹⁷ Note that it is important to the case that the cause of the self-doubt is purely psychological, as otherwise we would have grounds for thinking that the basis for belief is changing. This example is an adaption of Radford's (1966) 'diffident schoolboy' case, albeit to illustrate a very different point.

¹⁸ There may be some versions of the epistemic basis relation that wouldn't allow this, such as Swain's (1981) counterfactual account. But I think that this would be a problem for such a view, and hence, if true, would demand revision of the proposal. For more on the epistemic basing relation, see Neta

That same basis could just as much be a basis for *disbelief* in this proposition as belief, for example. But if that's right, then the presence of the basis must also be compatible with not forming a belief in the target proposition, given that disbelief entails, but is not entailed by, a lack of belief in the proposition.

We can evade this problem if we shift to an anti-risk epistemology. For our concern now is no longer with the modal closeness of the non-obtaining of the event in question (forming a true belief), but rather with the modal closeness of the risk event. When it comes to knowledge, I think the risk event is very clear – it's that the very same basis for belief leads you into error (i.e. false belief). That is, we want a safe basis for belief that will ensure that we can't very easily end up in error. So, this problem disappears according to an anti-risk epistemology, as we now have a principled way of focusing on the modal closeness of the basis leading to false belief, specifically, rather than either false belief or non-belief.

In fact, the move to anti-risk epistemology also better motivates the way that safety-based theorists like myself respond to the problem posed by necessary propositions.¹⁹ The general contours of the puzzle is that one could form a belief in a necessary proposition in a completely haphazard way that has no epistemic credentials at all – guesswork, say – but that it would inevitably be safe nonetheless because there is no close possible world where one can believe such a proposition and believe it falsely (since there is no possible world where it is false). But this appears wrong, in that we would want to say that any belief formed in this way is unsafe; it just seems a matter of luck that one happened on the right answer.²⁰ The way safety theorists like myself respond to this is to say that once we shift to a basis-relative formulation of safety – which, recall, was independently motivated in terms of an anti-luck epistemology – then our focus should not be on the particular proposition believed in the actual world, but

2011 and Korcz 2015. See also Bondy and Pritchard 2016, which makes the case for a new kind of epistemic luck regarding the basing relation, and motivates the proposal via appeal to anti-risk epistemology.

¹⁹ In fact, they don't have to be necessary for the problem to arise. So long as the truth of the proposition is modally stable across close possible worlds, then the same problem will arise, even if the proposition in question is contingent.

²⁰ It would be different if the world were *engineered* to guarantee you true beliefs within a certain domain (e.g. by a helpful demon). Now I would grant that your beliefs are safe, though they still aren't knowledge. Indeed, the 'temp' case that I offer as part of a critique of (what I call) 'pure' (or 'robust') anti-luck epistemology (i.e. an account of knowledge that only has an anti-luck epistemic condition) is meant to illustrate that one can have beliefs that are bound to be true (and hence are not luckily true) but which do not amount to knowledge. For discussion, see Pritchard, Millar, and Haddock 2010: ch. 3 and Pritchard 2012a.

rather on the doxastic output of the basis in the actual world instead. This means that while there is, of course, no close possible world where one falsely believes the necessary proposition that one actually believes, the kind of haphazard basis described above will lead to lots of false beliefs in close possible worlds (just not false beliefs in the proposition actually believed). As such it will be an unsafe basis.²¹

This certainly fixes the problem. But the motivation for this formulation of safety from an anti-luck epistemology point of view is rather weak. Yes, this licenses a focus on a basis-relative account of safety, but notice that one could have such an account and nonetheless insist on it being a basis for belief in the proposition actually believed, so there is still a motivational lacuna in play here. In particular, what aspect of anti-luck epistemology is meant to motivate the loosening of this restriction so that our focus should be on the doxastic output of the actual basis more generally in close possible worlds, and not simply on whether that basis leads to a false belief in close possible worlds regarding the proposition actually believed?

This lacuna disappears once we adopt anti-risk epistemology. For notice that we have explicitly formulated the risk event in terms of the basis leading us into error (false belief). If that's our focus, however, then of course we will not want to restrict our attentions only to error apropos the proposition actually believed, for then we will miss out on the wider errors that a bad basis can lead us to, as when a haphazard approach to belief-formation leads to lots of false beliefs in close possible worlds.

Finally, notice that anti-risk epistemology enables a much wider range of epistemic assessments than anti-luck epistemology. This is because, as noted above, we can introduce a range of different risk events when making a risk assessment, including considering multiple risk events within a single assessment. When it comes to knowledge, as I just noted, it seems clear that there is a single risk event that concerns us, which is that our basis for belief will not lead us into error (false belief) – that's the risk event that we want to be modally distant in order for one to count as a knower.

But once we turn our attentions away from knowledge, specifically, then we might have other risk events in mind, including multiple risk events. Think about good inquiries, for example. Sure, ending up with a false belief via a particular basis might well be one of the core risk events to avoid in a good inquiry, but it needn't be the full story. It could be, for example, that the good inquirer also needs to come up with an answer

²¹ I develop this line in a number of places. See, for example, Pritchard 2012a, 2012b. See also Pritchard 2009b.

within a reasonable timeframe, such that another risk event, to weigh-up against the former, is that one altogether fails to resolve one's inquiry. Sometimes a suboptimal solution is better than no solution at all, after all. One could regard the good inquirer as trading these two risk events off against each other, and so forming a judgement about what the right lines of inquiry should be. That is, one might be willing to be sanguine about a higher level of risk that one's inquiry leads to error than is normal in order to be more confident that one will avoid the other risk scenario of failing to resolve the inquiry.

This would, of course, add a pragmatic element into the nature of good inquiry, but I take it that this isn't particularly controversial, unlike a pragmatic encroachment thesis in the theory of knowledge.²² On a traditional account of the nature of the virtues, after all, the intellectual virtues are entwined with the practical and moral virtues, and so one would not expect the manifestation of the former to be in isolation from the latter.²³ More precisely, while a particular intellectually virtuous line of inquiry, once undertaken, might be uninfluenced by purely pragmatic factors, the question of which line of inquiry to undertake, and relatedly the criteria under which that inquiry is to be evaluated, may well be highly influenced by practical factors. In this way, one's intellectual virtues can work in concert with one's other virtues to promote one's (overall) flourishing. Similarly, the question of which risk event or events are relevant to conducting an inquiry might be a purely practical consideration, even if the inquiry itself is undertaken in a purely intellectual spirit. Pragmatic encroachment about inquiry of this kind is nothing to be concerned about.

Note too that the use of multiple risk events is especially salient once we think about inquiries as collaborative endeavours (as they often are). Consider a particular realm of scientific inquiry. It might be overall epistemically very beneficial that within that realm there are scientists taking varying degrees of epistemic risk, with some being very risk-adverse

²² For more on pragmatic encroachment in the theory of knowledge, see Hawthorne 2004; Stanley 2005; and Fantl and McGrath 2012. Note that Ballantyne 2011, 2012 has argued that anti-luck epistemology is committed to pragmatic encroachment about knowledge – a claim that, if sound, would plausibly carry over to anti-risk epistemology. Ballantyne's argument, however, depends on an earlier version of the modal account of luck that I offered which included a significance condition (e.g., Pritchard 2005). This extra condition is now dropped, for reasons that I articulate in Pritchard 2014. Similarly, there is no significance condition in the modal account of risk that I offer either, and hence there is no basis (in this particular regard anyway) for thinking that either anti-luck or anti-risk epistemology is committed to a pragmatic encroachment thesis about knowledge.

²³ For an influential neo-Aristotelian account of the virtues along these lines, with a particular focus on the nature of the intellectual virtues and their role within the virtues more generally, see Zagzebski 1996.

and so pursuing relatively safe lines of inquiry (such that they are likely to get the results they want, but those results are rarely going to be significant or surprising), and others taking large risks with very speculative lines of inquiry, with the potential of discovering important scientific truths (but more often than not discovering nothing of the kind). We can make sense of this very easily within an anti-risk epistemology, since what is taking place here is effectively a trade-off when it comes to the relative weights of different epistemic risk events.

This last point demonstrates that while anti-luck epistemology delivers a theory of knowledge (in anti-luck virtue epistemology), anti-risk epistemology goes one stage further and offers us not just a theory of knowledge (anti-risk virtue epistemology, as explored in Section 9.4), but also a broader way of approaching epistemic issues. In particular, we now have a richer way of evaluating epistemic standings in terms of a range of epistemic risk events, including even the possibility of multiple epistemic risk events, or even epistemic risk events considered from a specifically social perspective.

9.4 Anti-risk Virtue Epistemology

I noted above that I have argued elsewhere that anti-luck epistemology generates a theory of knowledge that I have termed anti-luck virtue epistemology. Anti-luck epistemology by itself cannot offer a complete theory of knowledge but can only deliver us the anti-luck condition on knowledge. In short, the reason for this is that such a condition is only going to give you a certain modal profile for one's belief. But that one's belief has this modal profile does not suffice to demonstrate that your cognitive success (i.e. true belief) is in any significant way down to your manifestation of relevant cognitive agency, something that I take to be integral to our concept of knowledge, as illustrated by the ability platitude noted above. In particular, we can imagine cases where external factors engineer it to be the case that there is no luck involved in one having true beliefs, and so one's beliefs are safe, but since one's cognitive success has nothing to do with one's manifestation of relevant cognitive agency, then one wouldn't count as a knower. This means that there needs to be an ability, or virtue, condition on knowledge, a condition that demands that there is the salient explanatory link between one's cognitive success and one's manifestation of relevant cognitive agency.²⁴

²⁴ For more on this point, see Pritchard, Millar, and Haddock 2010: ch. 3 and Pritchard 2012a.

Going in the other direction, I've also argued that one cannot account for knowledge exclusively in terms of a virtue-theoretic condition. That is, one can't simply subsume anti-luck epistemology within a broader virtue-theoretic account of knowledge, as a number of virtue theorists have tried to do (I call them *robust virtue epistemologists*). Very roughly, robust virtue epistemologists hold that knowledge is cognitive success that is *because of* the manifestation of relevant cognitive agency, where this means that the former is primarily attributable to the latter.²⁵ The thought is that while Gettier-style cases clearly do involve virtuous agency and cognitive success, they don't enjoy the relational quality of being a cognitive success that is because of one's virtuous agency. In short, they are not *cognitive achievements*, where this means a cognitive success that is because of one's manifestation of relevant cognitive agency (and which is a subclass of the broader category of achievements *simpliciter* – successes that are because of one's manifestation of relevant agency). In any case, the claim is that once one adds this relational ingredient into the mix then virtue epistemology can deal with the problem posed by epistemic luck without needing to appeal to anti-luck epistemology. I maintain that this is mistaken.

I originally argued for this point by appealing to a distinction between the kind of *intervening epistemic luck* that is familiar from standard Gettier-style cases, and the *environmental epistemic luck* that one finds in specifically barn-facade-type cases.²⁶ In standard Gettier-style cases something intervenes between one's rational basis for the target proposition and the fact itself, though one's belief is true nonetheless. So, for example, one reasonably thinks that one is seeing a sheep in the field, and there is a sheep in the field, but in fact what one is looking at is a big hairy dog that is obscuring from view the real sheep behind.²⁷ Cases of environmental epistemic luck are different, however, in that nothing intervenes in this

²⁵ This is how Greco 2003, 2007b, 2008, 2009a, 2009c understands the 'because of' relation, though see Greco 2012 for a reworking of his view. The other main proposal in the literature in this regard is due to Sosa 1991, 2007, 2009, 2011, 2015. This construes the 'because of' relation in terms of disposition manifestation. That a glass is shattered when hit, for example, could be because it is fragile, where this kind of explanation need not be in competition with a causal explanatory story (e.g. that so-and-so lost his temper and threw the glass at the wall). For specific discussion of Sosa's account, see Pritchard 2009a and Kallestrup and Pritchard 2016. See also Zagzebski 1996, 1999, who treats the 'because of' relation as an indefinable primitive. In order to keep the discussion to a manageable length, I will be focusing on the causal-explanatory construal of the 'because of' relation.

²⁶ I introduced this distinction, and the associated terminology, in Pritchard 2009a, 2009b, 2009c. I have further discussed this distinction in a number of places, including Pritchard, Millar, and Haddock 2010: chs 2–4, Pritchard 2012a, 2015a, and Kallestrup and Pritchard 2014.

²⁷ This is Chisholm's (1977: 105) famous Gettier-style case.

way. In the barn facade case, for example, one really does see a genuine barn, it is just that one is in an environment that ensures that one's belief is nonetheless unsafe.

Here is the thing: environmental luck (epistemic or otherwise) is entirely compatible with one's success (cognitive or otherwise) being primarily attributable to the manifestation of one's relevant (cognitive or otherwise) agency. Consider shooting an arrow at a target. If there is merely intervening luck involved, such that you skilfully took the shot, and was successful, but the latter was just down to an intervention (e.g. a dog running on, grabbing the bolt, and putting it in the target), then this success is not primarily attributable to the manifestation of your agency. It is not, as we say, your achievement at all, even though you performed well.

But now consider a parallel case of environmental luck. Everything is the same except this time the dog failed at the last moment to intercept the bolt. In close possible worlds the dog would have intercepted it, it's just that he failed in the actual world. Your success is modally fragile, and thus unsafe. Nonetheless, isn't this success now primarily down to your manifestation of relevant cognitive agency? After all, while something could have intervened – this is what ensures that the success is unsafe – in fact it didn't. This means that environmental luck (epistemic or otherwise) is compatible with genuine achievements. It also means, in turn, that achievements (cognitive or otherwise) can be unsafe. But since knowledge demands safety, this entails that we cannot equate knowledge with cognitive achievements, as the proponent of robust virtue epistemology proposes.

These days I make essentially the same point in terms of the notion of an *epistemic twin earth* case, which is effectively a way of sharpening up the earlier distinction between intervening and environmental epistemic luck.²⁸ Imagine two counterpart agents, one on earth and one on twin earth, who are microphysical duplicates of each other, with identical causal histories. The environment that they are presently causally interacting with is also identical in every respect. In addition, the 'normal' environment – i.e. the sorts of things that they would normally be causally interacting with – is also identical for both subjects. Now suppose that both agents form, on the same basis, the belief that *p*. All that is different with regards to the two agents is that the agent on twin earth occupies a very different modal environment. For whereas the agent on earth is forming a belief that *p* in such a way that this same basis for belief will generate true beliefs across all close possible worlds, due to idiosyncratic features of their modal

²⁸ See Kallestrup and Pritchard 2014. See also Pritchard 2016a.

environmental the agent on twin earth is forming a belief that p in such a way that the very same basis for belief will generate false beliefs in close possible worlds.

What is interesting about epistemic twin earth cases is that we have effectively kept fixed across the two subjects any possible factor that might be relevant to the manifestation of cognitive agency. After all, one can imagine that manifestations of agency – cognitive or otherwise – might be influenced by such factors as one's actual causal environment, one's normal causal environment, one's causal history, or one's microphysical nature. But no-one holds that manifestations of agency – again, cognitive or otherwise – are influenced by factors that are exclusive to one's modal environment. And yet the agent on twin earth is forming their belief in such a way that it is unsafe, in contrast to the agent on earth who is forming an identical belief. Insofar as we grant that knowledge is incompatible with veritic epistemic luck (and hence cannot be unsafe), then we should be inclined to treat the agent's belief on twin earth as not being knowledge.²⁹

I take the forgoing to indicate that the anti-luck and ability platitudes about knowledge, noted above, in fact impose distinct constraints on a theory of knowledge, rather than, as is (implicitly) supposed, the same constraints. One can see why they might be thought to impose the same constraints. After all, in general at least, the reason why one's true belief is due to luck (and hence fails to satisfy the anti-luck constraint) is that one's true belief had nothing to do with one's manifestation of relevant cognitive agency. Going in the other direction, the usual reason why one's true belief was not because of one's manifestation of relevant cognitive agency is that it was due to luck.

But this *prima facie* account of how the two intuitions relate has been shown to be problematic. On the one hand, we find that there are cases where agents have beliefs that have the right modal profile to count as non-lucky (safe), but which don't satisfy the ability constraint on knowledge. On the other hand, epistemic twin earth cases, and the phenomenon of environmental epistemic luck more generally, show that an agent's cognitive success can be primarily creditable to their manifestation of cognitive agency and yet nonetheless be lucky (unsafe). The upshot is that we need an account of knowledge that recognises the way in which these overlapping constraints on one's theory sometimes diverge.

²⁹ For further discussion of cases involving purely modal veritic epistemic luck, see Pritchard 2015a: §3.

This is where anti-luck virtue epistemology came in. Roughly, this argues that knowledge is safe cognitive success that is significantly attributable to one's manifestation of relevant cognitive agency. Note that this is both in a sense stronger and in a sense weaker than robust virtue epistemology. The former, because the explanatory relation is now between *safe cognitive success* and the manifestation of relevant cognitive agency rather than just *cognitive success*. The latter, because we are now only demanding that this explanatory relation satisfy a *significant* level rather than the cognitive success being primarily attributable to the cognitive agency of the subject.³⁰ Such a proposal can capture the way our two core platitudes intersect with one another. Sometimes even a high manifestation of cognitive agency, of a kind that would ordinarily easily suffice for knowledge, can fail to bring one knowledge due to purely environmental factors that make the cognitive success that results nonetheless unsafe. And sometimes a relatively low manifestation of cognitive agency, of a kind that wouldn't ordinarily deliver knowledge, can suffice for knowledge on account of the fact that one is in the kind of epistemically friendly environment that ensures that the belief is safe nonetheless. As I've argued elsewhere, both kinds of scenarios are uniquely accommodated by an anti-luck virtue epistemology.³¹

Crucially, anti-risk virtue epistemology, which is just anti-luck virtue epistemology with an anti-risk construal of safety, will inherit all the benefits of anti-luck virtue epistemology, while facing none of the problems that we saw afflicted the motivation of safety on anti-luck grounds. In particular, the kinds of theoretical lacunae that we noted as regards anti-luck epistemology will not infect anti-risk virtue epistemology. But this is not the only advantage to replacing an anti-luck virtue epistemology with an anti-risk virtue epistemology.

In order to see this, we need to note that, unsurprisingly, the proponents of robust virtue epistemology are not convinced that they should abandon their view and endorse such a two-aspect proposal. As noted

³⁰ Robust virtue epistemology needed the stronger claim in this regard in order to deal with Gettier-style cases, but we don't face that hurdle, as of course we have safety built into the proposal from the off. This part of the proposal helps us to accommodate cases of knowledge where one's cognitive success is primarily down to the cognitive agency of others, as is quite common in testimonial cases. See, for example, Pritchard 2012a for discussion of both points.

³¹ In particular, I claim that these two scenarios are both cases of *epistemic dependence*, where what I mean by that is how knowledge can be significantly influenced by factors outwith one's cognitive agency. The first kind of scenario is *negative epistemic dependence*, whereby environmental factors prevent an otherwise high manifestation of cognitive agency from being knowledge. The second kind of scenario is *positive epistemic dependence*, whereby environmental factors enable an otherwise low manifestation of cognitive agency to count as knowledge. See, for example, Pritchard 2016a for the details.

previously, epistemic twin earth cases bring the phenomenon of environmental epistemic luck, and the problem it poses for robust virtue epistemology, into sharp relief. Such cases effectively make it impossible to evade the problem posed by environmental epistemic luck to robust virtue epistemology. Interestingly, however, the response from robust virtue epistemologists on this score is increasingly not to try to evade the problem at all, but rather to meet it head-on by arguing that one can coherently allow that knowledge is compatible with environmental epistemic luck. Accordingly, they argue that knowledge can be compatible with veritic epistemic luck after all (of this particular variety anyway).³²

You can see why they might take this line. After all, their concern is with cognitive achievements – i.e. cognitive successes that are because of cognitive agency – so if environmental epistemic luck doesn't undermine cognitive achievements, then why care about it? Why not insist that knowledge is a cognitive achievement, and hence is compatible with environmental epistemic luck? Relatedly, surely everyone will agree that as veritic epistemic luck goes intervening epistemic luck is much worse, from an epistemic point of view, than environmental epistemic luck. After all, in the latter case, there is no mismatch between one's reasons and the facts that one finds in intervening epistemic luck.

This is where the shift to anti-risk epistemology, and thus to an anti-risk virtue epistemology, becomes important. One might be sanguine about one's theory of epistemology ascribing knowledge in cases where there is a high level of epistemic luck. But can one really be sanguine about one's theory of knowledge ascribing knowledge in cases where there is a high level of epistemic risk? Cases of environmental epistemic luck involve levels of epistemic risk that are on a par with cases of intervening epistemology luck, after all, and even robust virtue epistemologists are inclined to treat the latter as problematic from the perspective of ascribing knowledge. So how can they now consistently argue that we can be sanguine about one form of high epistemic risk if they themselves agree that corresponding levels of risk are elsewhere incompatible with knowledge?

Moreover, recall our point that we tend to care about (bad) luck because we care about risk, rather than vice versa. This means, in the epistemic case, that we tend to worry about epistemic luck because we are worried about epistemic risk, rather than vice versa. Accordingly, it is not the high

³² See, in particular, Sosa's (2007: ch. 5) discussion of the 'jokester' case, which has been very influential. I critically discuss Sosa's reasons for ascribing knowledge in this case in Pritchard 2009a.

levels of epistemic luck, specifically, that we should focus upon, but rather the high levels of epistemic risk. With the attentions of robust virtue epistemologists duly focused on the latter, it is hard to see how they could consistently argue that they are content to allow knowledge to be compatible with this high level of epistemic risk, particularly when elsewhere they maintain that such a level of epistemic risk is incompatible with knowledge. The upshot is that once anti-risk epistemology, and hence anti-risk virtue epistemology, is on the table, then it is far harder for proponents of robust virtue epistemology to shrug-off the high levels of epistemic risk involved in cases of environmental epistemic luck.³³

Anti-risk epistemology also gives us a very specific handle on how to think about post-Gettier epistemology. Recall that we noted above that luck is essentially a backward-facing notion, in contrast to risk, which is by its nature a forwards-looking notion. This point is important to diagnosing the direction post-Gettier epistemology took. For notice that Gettier-style cases are by their nature examples whereby the agent was cognitively successful, and we are reflecting back on the unusual, happenstance, way in which they were cognitively successful. It is thus no surprise that our focus becomes on the role of luck in this cognitive success, and thus on epistemic luck.

But if I am right that our interest in eliminating epistemic luck from knowledge is derivative on our interest in eliminating (high levels of) epistemic risk, then this backwards-focused nature of Gettier-style assessments led us to focus on quite the wrong thing, particularly given the (subtle, yet important) differences between luck and risk, and hence between epistemic luck and epistemic risk. In particular, had we spotted this distinction earlier, and clocked the need to examine these cases from the forwards-looking perspective provided by risk, then we would have been better placed to resolve the problems in hand (i.e. not just the narrow Gettier problem of coming up with an anti-luck/risk condition on knowledge, but also the broader Gettier problem of developing an adequate theory of knowledge).³⁴

9.5 Conclusion

We have seen that the shift from anti-luck epistemology to anti-risk epistemology, while on the face of it a rather marginal change, in fact

³³ I first made this point in Pritchard 2016b. See also Pritchard 2017a.

³⁴ I distinguish between these two interpretations of the 'Gettier problem', and offer resolutions of these problems – albeit via an anti-luck, rather than an anti-risk, epistemology – in Pritchard 2015a, 2017b.

offers us important new theoretical resources, on account of the small – but crucial nonetheless – differences between the notions of luck and risk. Anti-risk epistemology better motivates certain epistemic moves that we need to make, thereby removing significant motivational lacunae that were present in anti-luck epistemology. It also offers a new approach to epistemic questions that allows us to broaden our scope of epistemic evaluations by focusing upon a range of different, and potentially multiple, epistemic risk events. But perhaps most importantly, we have seen that this change of direction enables us to offer a more compelling virtue-theoretic account of knowledge. As we have seen, anti-risk virtue epistemology inherits all the advantages of anti-luck virtue epistemology but none of its flaws. Crucially, it presents us with the means to better resist the current fashion of taking a more permissive line on unsafe beliefs by illustrating how this commits one to implausible claims about epistemic risk, while at the same time providing us with a better diagnostic handle on why post-Gettier epistemology panned in out in quite the manner that it did. So, while the shift from anti-luck epistemology to anti-risk epistemology may seem of minor philosophical significance at first blush, on close inspection we see that it delivers important theoretical gains.³⁵

REFERENCES

- Ballantyne, N. 2011. 'Anti-luck Epistemology, Pragmatic Encroachment, and True Belief', *Canadian Journal of Philosophy* 41: 485–504.
 2012. 'Luck and Interests', *Synthese* 185: 319–334.
 Bondy, P., and Pritchard, D. H. 2016. 'Propositional Epistemic Luck, Epistemic Risk, and Epistemic Justification', *Synthese* (Online first, DOI: 10.1007/s11229-016-1262-2).
 Broncano-Berrocal, F. 2015. 'Luck as Risk and the Lack of Control Account of Luck', *Metaphilosophy* 46: 1–25.
 Chisholm, R. 1977. *Theory of Knowledge* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall.
 Coffman, E. J. 2015. *Luck: Its Nature and Significance for Human Knowledge and Agency*. London: Palgrave Macmillan.
 Fantl, J., and McGrath, M. 2012. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

³⁵ Thanks to Jesus Navarro for helpful discussions on this topic, and to John Greco and Chris Kelp for detailed comments on an earlier draft. An earlier version of this chapter was presented at a workshop on contemporary issues in epistemology at the University of California, Irvine in March 2018, and at a conference on 'New Trends in Virtue Epistemology' at the University of Glasgow in September 2018. I am grateful to the audience for their questions on both occasions, and especially Michael Brady, who was the commentator on my paper at the Glasgow conference.

- Greco, J. 2003. 'Knowledge as Credit for True Belief', in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, M. DePaul and L. Zagzebski (eds). Oxford: Oxford University Press, 111–134.
- 2007a. 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17: 57–69.
- 2007b. 'Worries about Pritchard's Safety', *Synthese* 158: 299–302.
2008. 'What's Wrong with Contextualism?', *Philosophical Quarterly* 58: 416–436.
- 2009a. *Achieving Knowledge*. Cambridge, UK: Cambridge University Press.
- 2009b. 'Knowledge and Success from Ability', *Philosophical Studies* 142: 17–26.
- 2009c. 'The Value Problem', in *Epistemic Value*, A. Haddock, A. Millar, and D. H. Pritchard (eds). Oxford: Oxford University Press, 313–321.
2012. 'A (Different) Virtue Epistemology', *Philosophy and Phenomenological Research* 85: 1–26.
- Hansson, S. O. 2004. 'Philosophical Perspectives on Risk', *Techné: Research in Philosophy and Technology* 8: 10–35.
2014. 'Risk', *Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.). <http://plato.stanford.edu/archives/spr2014/entries/risk/>.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Kahneman, D., and Varey, C. A. 1990. 'Propensities and Counterfactuals: The Loser That Almost Won', *Journal of Personality and Social Psychology* 59: 1101–1110.
- Kallestrup, J. and Pritchard, D. H. 2014. 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy* 22: 335–357.
- Korcz, K. A. 2015. 'The Epistemic Basing Relation', *Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.). <https://plato.stanford.edu/entries/basing-epistemic/>.
- Lackey, J. 2008. 'What Luck Is Not', *Australasian Journal of Philosophy* 86: 255–267.
- Levy, N. 2011. *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford: Oxford University Press.
- Neta, R. 2011. 'The Basing Relation', *The Routledge Companion to Epistemology*, S. Bernecker and D. H. Pritchard (eds). New York: Routledge, 109–118.
- Pritchard, D. H. 2004. 'Epistemic Luck', *Journal of Philosophical Research* 29: 193–222.
2005. *Epistemic Luck*, Oxford: Oxford University Press.
2007. 'Anti-luck Epistemology', *Synthese* 158: 277–297.
2008. 'Sensitivity, Safety, and Anti-luck Epistemology', in *Oxford Handbook of Scepticism*, J. Greco (ed.). Oxford: Oxford University Press, 437–455.
- 2009a. 'Apt Performance and Epistemic Value', *Philosophical Studies* 143: 407–416.
- 2009b. *Knowledge*. London: Palgrave Macmillan.
- 2009c. 'Knowledge, Understanding and Epistemic Value', *Epistemology (Royal Institute of Philosophy Lectures)*, A. O'Hear (ed.). Cambridge, UK: Cambridge University Press, 19–43.

- 2012a. 'Anti-luck Virtue Epistemology', *Journal of Philosophy* 109: 247–279.
- 2012b. 'In Defence of Modest Anti-luck Epistemology', *The Sensitivity Principle in Epistemology*, T. Black and K. Becker (eds). Cambridge, UK: Cambridge University Press, 173–192.
2014. 'The Modal Account of Luck', *Metaphilosophy* 45: 594–619.
- 2015a. 'Anti-luck Epistemology and the Gettier Problem', *Philosophical Studies* 172: 93–111.
- 2015b. 'Risk', *Metaphilosophy* 46: 436–461.
- 2016a. 'Epistemic Dependence', *Philosophical Perspectives* 30: 1–20.
- 2016b. 'Epistemic Risk', *Journal of Philosophy* 113: 550–571.
- 2017a. 'Anti-risk Epistemology and Negative Epistemic Dependence', *Synthese* (Online first: <https://doi.org/10.1007/s11229-017-1586-6>).
- 2017b. 'Knowledge, Luck and Virtue: Resolving the Gettier Problem', in *Explaining Knowledge: New Essays on The Gettier Problem*, C. Almeida, P. Klein, and R. Borges (eds). Oxford: Oxford University Press, 57–73.
- 2017c. 'Legal Risk, Legal Evidence, and the Arithmetic of Criminal Justice', *Jurisprudence* (Online first, DOI: 10.1080/20403313.2017.1352323).
2018. 'Aesthetic Risk', *Think* 17: 1–14.
- Pritchard, D. H., Millar, A., and Haddock, A. 2010. *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press.
- Pritchard, D. H., and Smith, M. 2004. 'The Psychology and Philosophy of Luck', *New Ideas in Psychology* 22: 1–28.
- Pritchard, D. H., and Whittington, L. J. (eds). 2015. *The Philosophy of Luck*. Oxford: Blackwell.
- Radford, C. 1966. 'Knowledge – By Examples', *Analysis* 27: 1–11.
- Riggs, W. 2007. 'Why Epistemologists Are so down on Their Luck', *Synthese* 158: 329–344.
2009. 'Luck, Knowledge, and Control', *Epistemic Value*, A. Haddock, A. Millar, and D. H. Pritchard (eds). Oxford: Oxford University Press, 205–221.
- Sainsbury, R. M. 1997. 'Easy Possibilities', *Philosophy and Phenomenological Research* 57: 907–919.
- Sosa, E. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge, UK: Cambridge University Press.
1999. 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13: 141–154.
2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
2011. *Knowing Full Well*. Princeton, NJ: Princeton University Press.
2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Swain, M. 1981. *Reasons and Knowledge*. Ithaca, NY: Cornell University Press.

- Teigen, K. H. 1995. 'How Good Is Good Luck?: The Role of Counterfactual Thinking in the Perception of Lucky and Unlucky Events', *European Journal of Social Psychology* 25: 281–302.
1996. 'Luck: The Art of a Near Miss', *Scandinavian Journal of Psychology* 37: 156–171.
1997. 'Luck, Envy, Gratitude: It Could Have Been Different', *Scandinavian Journal of Psychology* 38: 318–323.
- 1998a. 'Hazards Mean Luck: Counterfactual Thinking and Perceptions of Good and Bad Fortune in Reports of Dangerous Situations and Careless Behaviour', *Scandinavian Journal of Psychology* 39: 235–248.
- 1998b. 'When the Unreal is More Likely Than the Real: *Post Hoc* Probability Judgements and Counterfactual Closeness', *Thinking and Reasoning* 4: 147–177.
2003. 'When a Small Difference Makes a Large Difference: Counterfactual Thinking and Luck', in *The Psychology of Counterfactual Thinking*, D. R. Mandel, D. Hilton, and P. Catellani (eds). London: Routledge, 129–146.
- Tetlock, P. E. 1998. 'Close-Call Counterfactuals and Belief-System Defenses: I Was Not Almost Wrong but I Was Almost Right', *Journal of Personality and Social Psychology* 75: 639–652.
- Tetlock, P. E. and Lebow, R. N. 2001. 'Poking Counterfactual Holes in Covering Laws: Cognitive Styles and Historical Reasoning', *American Political Science Review* 95: 829–843.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge, UK: Cambridge University Press.
1999. 'What Is Knowledge?', *Blackwell Guide to Epistemology*. J. Greco and E. Sosa (eds). Oxford: Blackwell, 92–116.