THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Epistemic risk

Citation for published version:

Pritchard, D 2016, 'Epistemic risk', *Journal of Philosophy*, vol. 113, no. 11, pp. 550-571.
<https://doi.org/10.5840/jphil20161131137>

Digital Object Identifier (DOI):

[10.5840/jphil20161131137](https://doi.org/10.5840/jphil20161131137)

Link:

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

Document Version:

Peer reviewed version

Published In:

Journal of Philosophy

Publisher Rights Statement:

This is the peer-reviewed version of the following article: Pritchard, D. (2016). Epistemic risk. *The Journal of Philosophy*, 113, 550–571., which has been published in final form at [dx.doi.org/10.5840/jphil20161131137](https://doi.org/10.5840/jphil20161131137). This article may be used for non-commercial purposes only.

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.

EPISTEMIC RISK

The goal of this paper is to mark the transition from an *anti-luck epistemology* to an *anti-risk epistemology*, and to explain in the process how the latter has advantages over the former. We begin with an account of anti-luck epistemology and the *modal account of luck* that underpins it. Then we will consider the close inter-relationships between luck and risk, and in the process set out the *modal account of risk* that is a natural extension of the modal account of luck. Finally, we will apply the modal account of risk to epistemology in order to develop an anti-risk epistemology, and then explore the merits of this proposal. In particular, it is shown that (i) this account can avoid a theoretical lacuna in anti-luck epistemology, and (ii) there is a stronger theoretical motivation for anti-risk epistemology compared with anti-luck epistemology, especially when it comes to explaining why *environmental epistemic luck* is incompatible with knowledge.

I. ANTI-LUCK EPISTEMOLOGY AND THE MODAL ACCOUNT OF LUCK IN OUTLINE

It is a widely held platitude in epistemology that knowledge is in some fundamental sense incompatible with luck. In particular, knowledge seems incompatible with the possibility that one's true belief is only true as a matter of luck, what is often known as *veritic epistemic luck*. Call this general intuition about the relationship between knowledge and luck the *anti-luck platitude*. Until quite recently this platitude was taken largely at face value, as something that did not require further elucidation. In contrast, the guiding thought behind *anti-luck epistemology* is precisely that one should *not* take the anti-luck platitude at face value but rather carefully unpack it. In particular, anti-luck epistemology urges a three-stage approach to the theory of knowledge that takes the anti-luck platitude as central to the project. First, one offers a theory of luck. Second, one delineates the specific sense in which knowledge is incompatible with luck. Finally, third, one puts these two component parts together and formulates an *anti-luck condition* on knowledge that captures the specific sense in which knowledge is incompatible with luck. If the anti-luck platitude does reveal something deep and important about knowledge, then by undertaking the anti-luck epistemological project one should determine a core epistemic condition on knowledge. Indeed,

one might even determine an epistemic condition that is, with true belief, sufficient, or close to being sufficient anyway, for knowledge. That would be quite a result.

One of the attractions of this theoretical project is that it might offer a principled way of dealing with the Gettier problem—*viz.*, the problem of explaining what it takes to avoid the specific kind of epistemic luck that undermines knowledge in Gettier-style cases. On standard epistemological proposals one tries to deal with the Gettier problem by working out which condition or conditions one needs to add to one's favoured non-Gettier-proof account of knowledge in order to make it Gettier-proof. This way of approaching the problem has been notoriously unsuccessful, and tends to lead to analyses of knowledge that strike one as *ad hoc* and unmotivated. Indeed, the general lack of success of this way of dealing with the Gettier problem has given rise to the widespread view that knowledge is not the kind of thing that is susceptible to an analysis.¹

Anti-luck epistemology promises to be a better way of approaching this issue since rather than focussing specifically on the kind of epistemic luck in play in Gettier cases, one instead attends to the general question of the nature of knowledge-undermining epistemic luck, whether it is found in Gettier-style cases or elsewhere. The anti-luck condition which results from a successfully conducted anti-luck epistemology will thus not be a mere anti-Gettier condition, even though it will exclude Gettier-style cases just as it excludes other cases of knowledge-undermining epistemic luck. Moreover, rather than taking the notion of luck as a primitive, anti-luck epistemology incorporates a theory of luck as a means of outlining the anti-luck condition on knowledge.

The theory of luck that I offer is what is known as the *modal account*. We can get a sense of the appeal of such a proposal by considering an insight from epistemology that I think is crucial for our understanding of luck. Consider the famous lottery example. Imagine a subject who holds a lottery ticket for a fair lottery with astronomically long odds, where the draw has been made. The ticket is a loser, but the subject has not yet heard the result and so has no inkling of this. We can add to this story in two interesting ways. In the first scenario, the subject becomes aware of the astronomical odds involved and hence on this basis alone forms the true belief that her ticket is a loser. In the second scenario, the subject hasn't paid any attention at all to the odds involved in

¹ Most famously, of course, this is the view defended by Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000). Note that we are here glossing over the issue of whether an adequate analysis of knowledge must thereby be a *reductive* analysis. My own view is that this is unnecessary, and that what we seek is rather an analysis that is *informative*. (Reductive analyses are sometimes uninformative, after all, as when they are *ad hoc*, and non-reductive analyses can nonetheless be informative). For more on the methodology of epistemology, see Duncan Pritchard, 'The Methodology of Epistemology', *Harvard Review of Philosophy* 18 (2012), 91-108, and 'Sceptical Intuitions', *Intuitions*, (eds.) D. Rowbottom & T. Booth, 213-31, (Oxford: Oxford University Press, 2014).

this lottery. Instead, she reads the result in a reliable newspaper and so on this basis forms the true belief that her ticket is a loser.

Here is the puzzle. It seems that the subject in the first scenario doesn't know that her ticket is a loser, and yet the subject in the second scenario does.² Moreover, the natural explanation of why this is so is that in the first scenario the subject's true belief is a matter of luck, while in the second scenario it is not a matter of luck. The reason why this is puzzling is that if we consider the subjects' bases for belief then, from a probabilistic point of view anyway, the odds in the second case are nothing like as massively in support of the truth of the subject's belief as they are in the first case. How then can it be that knowledge is present in the second case and not the first? Is knowledge not a straightforward function of the strength of one's evidence, probabilistically conceived?³

What the lottery case reminds us is that an event can be modally close even when probabilistically unlikely. That is, the possible world in which one wins a lottery, while probabilistically far-fetched, is in fact modally close. The possible world in which one is leaping about with joy in one's room because one is a lottery winner is very alike to the possible world in which one is tearing one's ticket up in disgust—all that needs to change is that a few coloured balls fall in a slightly different configuration.⁴

In contrast, the possible world in which a reliable newspaper misprints the lottery result, while not so probabilistically far-fetched, is not modally close. Newspapers have a morbid fear of printing erroneous lottery results—just think of the problems that this could cause—and so have elaborate systems in place in order to ensure the accuracy of what they print in this regard. It follows that one needs to change quite a lot about the actual world in order to get to the possible world where a reliable newspaper (*The Times*, say) prints an incorrect lottery result. The moral is that modal closeness comes apart from probabilistic closeness.⁵ In particular, one cannot infer from the fact that an event is probabilistically unlikely (such as a lottery win) that it is therefore also modally far-off.

² For empirical evidence that our folk judgements about lottery cases line-up with the relevant philosophical judgements, see John Turri and Ori Friedman, 'Winners and Losers in the Folk Epistemology of Lotteries', *Advances in Experimental Epistemology*, (ed.) J. Beebe, 45-70, (London: Continuum, 2014).

³ For more on the lottery problem, as it is known, see Pritchard, 'Knowledge, Luck, and Lotteries', *New Waves in Epistemology*, (eds.) D. H. Pritchard & V. Hendricks, 28-51, (London: Palgrave Macmillan, 2008). For a very different treatment of this problem, see John Hawthorne, *Knowledge and Lotteries*, (Oxford: Oxford University Press, 2014).

⁴ I am here characterizing possible worlds in the standard way in terms of a similarity ordering. See, for example, David Lewis, *Counterfactuals*, (Oxford: Blackwell, 1973), and Lewis, *On the Plurality of Worlds*, (Oxford: Blackwell, 1987). For further defence of this account of possible worlds in the context of the modal account of luck, see Pritchard, 'The Modal Account of Luck', *Metaphilosophy* 45 (2014), 594-619, §3.

⁵ As it happens, I have first-hand experience of this point about newspapers. In my late teens I gained work experience with a local newspaper and saw for myself the lengths they went to in order to ensure the accuracy of their lottery results. (And note that this is just a local newspaper with limited resources, rather than an internationally respected national newspaper like *The Times*).

Indeed, this is why people play lotteries, and yet do not place bets on modally far-fetched events with similarly massive odds. The odds of me winning the 100m gold medal at the next Olympics may well be in the region of your average lottery win, but you'd be crazy to bet the same amount you'd spend on a lottery ticket on this event obtaining. This is because not only is this event probabilistically far-fetched, but it is also modally far-fetched—an awful lot would need to change about the actual world to make it such that I am an Olympic sprint champion (indeed, I suspect it would take some sort of global conspiracy).

This distinction between modal and probabilistic closeness of events may seem highly theoretical, but it is nonetheless rooted in our ordinary judgements. The lottery example is a case in point, since we are surely sensitive to the fact that the lottery win scenario is modally close even while being probabilistically far-fetched (even if we wouldn't articulate this distinction in these terms of course). Indeed, that we recognise this point is revealed by our behaviour, in that we are not at all inclined to bet on modally far-fetched events with similar odds to the lottery.⁶

Why is knowledge lacking in the first scenario where the subject's true belief is based solely on the odds, but present in the second scenario where it is based on reading the result in a reliable newspaper? Here is a perfectly natural explanation. In the first scenario the subject's true belief is just down to luck, since she could so very easily have formed a false belief (i.e., had the balls fallen in a slightly different configuration, such that she owned the winning lottery ticket). As the point is sometimes put, given how the subject formed her belief, that her belief is false is an *easy possibility*.⁷ In the latter case, in contrast, the subject's true belief doesn't seem lucky at all. Given how she formed her true belief, she couldn't have easily formed a false belief, since reliable newspapers tend to publish the right result in worlds like the actual world.

Our judgements about knowledge are thus sensitive to the modal closeness of error as opposed to its probabilistic closeness. That is, the moral of the lottery example is that a true belief can fail to be knowledge even despite the odds being massively in its favour so long as the possibility of error is nonetheless modally close. In such cases we judge the agent's cognitive success to be too risky to count as knowledge.

This point about the distinction between modal and probabilistic closeness being rooted in

⁶ Note that the slogan for the UK's national lottery was, until quite recently, "It could be you!". This is clearly not the 'could' of probability, since in this sense it (realistically) couldn't be you, but rather the 'could' of modal nearness—i.e., if you play the lottery, then someone just like you will win it. This is borne out by their advertising campaign, which at one point featured a God-like finger hovering over ticket-holders, and then zapping one of them (the winner). Note that in arguing that one would be crazy to bet on a modally far-fetched event with similar odds to a lottery win I am not thereby suggesting that playing the lottery is rational. The point is rather that whatever one thinks of the rationality of playing the lottery, placing a bet on a modally far-fetched event with similar odds would be, from a rational point of view, much worse.

⁷ For more on the notion of an easy possibility, see R. M. Sainsbury, 'Easy Possibilities', *Philosophy and Phenomenological Research* 57 (1997), 907-19.

our everyday judgements is borne out by the empirical research on luck ascriptions. For example, in a series of studies conducted by the psychologist Karl Teigen it was found that when a success was perceived as being physically close to a failure (i.e., when a wheel of fortune stopped in a winning sector, but was physically close to stopping in a losing sector), the success was perceived as luckier than when the failure was not perceived as physically close. Moreover, he also found that this counterfactual closeness could not be understood simply in terms of the probabilities involved. Subjects were willing to treat events as being different as regards the degree of luck involved even whilst granting that the probabilities of each of the two events occurring was the same. Subjects would, for example, recognise that the probability of one's ball landing in a losing sector on a roulette wheel was constant wherever the ball landed in that losing sector, while also regarding an event in which one's ball landed near to the winning sector as involving bad luck, unlike other events where the ball landed further away (which, depending on where the ball landed, were either not regarded as unlucky at all, or else regarded as involving a lower degree of bad luck).⁸

With these kinds of considerations in mind, we are thus led to the modal account of luck. On this view, what makes an event lucky is that while it obtains in the actual world, there are—keeping the initial conditions for that event fixed—close possible worlds in which this event does not obtain. So, for example, a lottery win is a lucky event because there are close possible worlds where the initial conditions for this event are the same but where one does not win the lottery (i.e., where the coloured balls fall in a slightly different configuration).

What is meant here by the 'initial conditions for the event'? The point of this restriction is that we need to keep certain features of the actual world fixed in our evaluation of the close worlds. In particular cases, it is usually pretty clear what needs to remain fixed. In the lottery case, for example, we obviously need to keep fixed that the subject buys a lottery ticket and that the lottery retains much of its salient features (i.e., remains free and fair, with long odds, and so on). If, say, one were *guaranteed* to win the lottery (e.g., it is rigged in one's favour), then clearly this isn't a

⁸ See Karl Teigen, 'How Good is Good Luck?: The Role of Counterfactual Thinking in the Perception of Lucky and Unlucky Events', *European Journal of Social Psychology* 25 (1995), 281-302; 'Luck: The Art of a Near Miss', *Scandinavian Journal of Psychology* 37 (1996), 156-71; 'Luck, Envy, Gratitude: It Could Have Been Different', *Scandinavian Journal of Psychology* 38 (1997), 318-23; '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 (1998), 235-48; 'When the Unreal is More Likely Than the Real: *Post Hoc* Probability Judgements and Counterfactual Closeness', *Thinking and Reasoning* 4 (1998), 147-77; 'When a Small Difference Makes a Large Difference: Counterfactual Thinking and Luck', *The Psychology of Counterfactual Thinking*, (eds.) D. R. Mandel, D. Hilton & P. Catellani, 129-46, (London: Routledge, 2003). Other studies confirm Teigen's findings. See especially: Daniel Kahneman and C. A. Varey, 'Propensities and Counterfactuals: The Loser That Almost Won', *Journal of Personality and Social Psychology* 59 (1990), 1101-10; P. E. Tetlock, 'Close-Call Counterfactuals and Belief-System Defenses: I Was Not Almost Wrong but I Was Almost Right', *Journal of Personality and Social Psychology* 75 (1998), 639-5; and Tetlock and R. N. Lebow, 'Poking Counterfactual Holes in Covering Laws: Cognitive Styles and Historical Reasoning', *American Political Science Review* 95 (2001), 829-43. For a survey of recent psychological work on luck, see Pritchard and Matthew Smith, 'The Psychology and Philosophy of Luck', *New Ideas in Psychology* 22 (2004), 1-28.

lucky event even if, as it happens, there are close possible worlds in which one does not win the lottery (e.g., because one is prevented from buying a lottery ticket in such worlds).⁹

Is there a general specification that one can offer of these ‘initial conditions’? Well, we can say this much: they need to be specific enough to pick-out a particular kind of event that we want to assess for luckiness, but not so specific as to guarantee that this event obtains (e.g., we don’t want the purchase of a *winning* lottery ticket to be part of the initial conditions for the lottery win). That’s quite vague, of course, but my suspicion is that we shouldn’t expect anything more detailed, in that we shouldn’t require a theory to be any more precise than the phenomena about which we are theorizing. For our purposes it is enough that we can pick-out such initial conditions on a case-by-case basis (which I believe we usually can).

Note that this conception of luck can accommodate the idea that luck comes in degrees. Consider the lucky event of not being shot by a sniper’s bullet. With everything else kept fixed, imagine that in scenario A the bullet misses one by millimeters, whereas in scenario B the bullet misses one by a metre. We would naturally judge that both events are lucky, and our account of luck confirms this judgement, for in both cases there are close possible worlds in which the initial conditions for this event are the same and one is hit by the bullet. We would also naturally judge that scenario A is luckier than scenario B, and our account of luck again confirms this judgement. For the possible worlds in which one is hit by the bullet are clearly closer in scenario A than they are in scenario B, since less needs to be changed about the actual world in order to get to these possible worlds.

More generally, we can say that the degree of luck involved varies in line with the modal closeness of the world in which the target event doesn’t obtain (but where the initial conditions for that event are kept fixed). We would thus have a *continuum* picture of the luckiness of an event, from very lucky to not (or hardly) lucky at all. Once the degree of luck falls below a certain level—i.e., once there is no modally close world where the target event doesn’t obtain—then we would naturally classify the event as *not* lucky, since it does not involve a significant degree of luck. This conception of luck also allows us to compare events in terms of their luckiness, so even when considering two events which we don’t think are lucky we can nonetheless ask the question whether the one is luckier than the other (e.g., not being shot by a sniper is presumably luckier if the actual world is one where gun ownership is common than if, all other things being equal, gun ownership is rare—even when neither event involves a significant degree of luck).¹⁰

⁹ Of course, it may be lucky in such a case that one gets to buy a lottery ticket, but that’s a different event from the one under consideration, which is one’s winning of the lottery. (This point reminds us of the importance of making the target lucky event clear and keeping it fixed throughout our evaluation).

¹⁰ For discussion of how subjects’ judgements about degrees of luck vary in proportion to the counterfactual closeness of the target event, see Kahneman and Varey, *op. cit.*, and Teigen, ‘Luck: The Art of a Near Miss’, *op. cit.* For

Let us now return to anti-luck epistemology. We now have the theory of luck that we are looking for, in the form of the modal account of luck. In terms of the specific sense in which knowledge excludes luck, let us gloss over the dialectical twists and turns, of which there are many, and cut to the chase: we are interested in the event of the subject being cognitively successful (i.e., having a true belief), and we want this event to be non-lucky.¹¹ With our modal account in mind we can flesh this out by saying that a lucky cognitive success is a cognitive success where—keeping the relevant initial conditions fixed as usual—cognitive failure (i.e., false belief) is modally close. In such a case one's belief is subject to veritic epistemic luck, and hence doesn't amount to knowledge. A non-lucky cognitive success, in contrast, is thus one where, keeping the relevant initial conditions fixed, cognitive failure is not modally close. A bit more carefully, we can say that there is a continuum of epistemic luck in play here. Where the cognitive failure in question is modally very close, then the cognitive success is very lucky and hence knowledge is excluded. Where the cognitive failure is modally far-off, then the cognitive success is not significantly lucky and hence is compatible with knowledge. In between, there is spectrum of degrees of epistemic luck (and hence we would expect our judgements about knowledge possession to be more secure as the degree of epistemic luck lessens).

Anti-luck epistemology generates a number of interesting theoretical consequences. One of the overarching morals of anti-luck epistemology is that we should endorse the so-called *safety principle*, which is a modal condition on knowledge (and in particular, that we should endorse it over competing modal conditions on knowledge, such as the *sensitivity principle*).¹² Moreover, anti-luck epistemology motivates a particular rendering of this principle. The safety principle in outline demands a cognitive success that could not very easily have been a failure. So put, the general fit

a more detailed defence of the modal account of luck, including responses to the main objections and competing proposals, see Pritchard, 'The Modal Account of Luck', *op. cit.* For more on the philosophy of luck more generally, see Neil Levy, *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*, (Oxford: Oxford University Press, 2011), E. J. Coffman, *Luck: Its Nature and Significance for Human Knowledge and Agency*, (London: Palgrave Macmillan, 2015), and the papers collected in Pritchard and Lee Whittington, (eds.), *The Philosophy of Luck*, (Oxford: Wiley-Blackwell, 2015).

¹¹ For more on the details that we are eliding here—and in particular on the various kinds of epistemic luck that one can delineate, some of them malignant (i.e., knowledge-excluding) and some of them benign (i.e., knowledge-compatible)—see Pritchard, *Epistemic Luck*, (Oxford: Oxford University Press, 2005).

¹² For some key defences of (versions of) the safety principle (though not always under this description), see Steven Luper, 'The Epistemic Predicament', *Australasian Journal of Philosophy* 62 (1984), 26-50; Sainsbury, *op. cit.*; Ernest Sosa, 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13 (1999), 141-54; Williamson, *op. cit.*; and Pritchard, 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10 (2002), 283-307. For some key defences of the sensitivity principle, see Fred Dretske, 'Epistemic Operators', *Journal of Philosophy* 67 (1970), 1007-23, and 'Conclusive Reasons', *Australasian Journal of Philosophy* 49 (1971), 1-22; Robert Nozick, *Philosophical Explanations*, (Oxford: Oxford University Press, 1981); Sherrilyn Roush, *Tracking Truth: Knowledge, Evidence and Science*, (Oxford: Oxford University Press, 2005); Kelly Becker, *Epistemology Modalized*, (London: Routledge, 2007); Tim Black and Peter Murphy, 'In Defense of Sensitivity', *Synthese* 154 (2007), 53-71; and Black, 'Defending a Sensitive Neo-Moorean Invariantism', *New Waves in Epistemology*, (eds.) V. F. Hendricks & D. H. Pritchard, 8-27, (London: Palgrave Macmillan, 2008). For two recent comparative overviews of the literature on safety and sensitivity principles, see Pritchard, 'Sensitivity, Safety, and Anti-Luck Epistemology', *Oxford Handbook of Scepticism*, (ed.) J. Greco, 437-55, (Oxford: Oxford University Press, 2008), and Black, 'Modal and Anti-Luck Epistemology', *Routledge Companion to Epistemology*, (eds.) S. Bernecker & Pritchard, 187-98, (London: Routledge, 2011).

with our anti-luck epistemology is obvious.

But once we opt for an anti-luck epistemology we also get a very specific rendering of the safety principle. Rather than the general claim that a safe cognitive success is one that could not very easily have been a failure, we get instead a continuum picture of the epistemic luck involved, with modally very close epistemic luck incompatible with knowledge at the one end of the continuum, shading off along this continuum towards modally far-off epistemic luck that is compatible with knowledge. Moreover, given that we need to keep the initial conditions for the lucky event fixed, we get a basis-relative account of safety, whereby we keep fixed what gave rise to the original belief. The result is a much more nuanced conception of the safety principle.

As I've argued elsewhere, with the safety principle so understood we can deal with a range of problems that have been leveled against it.¹³ Furthermore, since this nuanced rendering of the safety principle is motivated on independent grounds, via a theory of luck, so there is nothing *ad hoc* about this objection-proof formulation of safety. It would take me too far afield to rehearse all the problems that have been posed for safety and outline how an anti-luck epistemology deals with them, but let me here at least describe one representative difficulty in this regard and how anti-luck epistemology resolves it.

One puzzle about the safety principle is why it doesn't entail that one's beliefs in necessary truths are automatically safe, even if they are formed in an epistemically lousy fashion. In order to see this, imagine that one forms one's belief that $2+2=4$ by tossing a coin. If one straightforwardly understands safety in terms of a safe belief that p , then the problem becomes immediately apparent, in that it obviously cannot be an easy possibility that one forms a false belief that $2+2=4$, since there is no possible world in which this proposition is false. And yet forming one's true beliefs, even as regards necessary truths, by flipping coins is surely veritically epistemically lucky in the relevant sense.

Anti-luck epistemology can explain what has gone wrong here. What we are interested in is whether the basis for belief in question could easily have led to the formation of a false belief (i.e.,

¹³ For example, I've argued that anti-luck epistemology offers a principled way of dealing with a dilemma that has been posed for safety-based theories of knowledge by John Greco, 'Worries About Pritchard's Safety', *Synthese* 158 (2007), 299-302. Moreover, I claim that anti-luck epistemology has predictive power, as it enables us to adjudicate, in a principled fashion, between opposing responses to certain cases, in that it can highlight which details of the case are important and thereby explain why two superficially identical formulations of a given example can generate very different responses. For more on anti-luck epistemology, see Pritchard, *Epistemic Luck*, *op. cit.*; 'Anti-Luck Epistemology', *Synthese* 158 (2007), 277-97; 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109 (2012), 247-79, and 'In Defence of Modest Anti-Luck Epistemology', *The Sensitivity Principle in Epistemology*, (eds.) T. Black & K. Becker, 173-92, (Cambridge: Cambridge University Press, 2012). See also Pritchard, 'Safety-Based Epistemology: Whither Now?', *Journal of Philosophical Research* 34 (2009), 33-45; 'There Cannot be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) M. Steup, J. Turri & E. Sosa, 152-64, (Oxford: Blackwell, 2013); 'Anti-Luck Epistemology and the Gettier Problem', *Philosophical Studies* 172 (2015), 93-111; and 'Knowledge, Luck and Virtue: Resolving the Gettier Problem', *The Gettier Problem*, (eds.) C. Almeida, P. Klein & R. Borges, (Oxford: Oxford University Press, *forthcoming*).

we are interested in whether the latter is an easy possibility, given that we keep the basis for belief fixed). Accordingly, what is relevant in cases where one forms one's belief in a necessary truth in an epistemically awful way (e.g., via a coin-toss) are not the close possible worlds where one forms a belief *in the very same proposition* and believes falsely (which is of course impossible), but rather the close possible worlds where one forms a belief on the same basis and believes falsely, where this belief might well be in a different proposition (e.g., the coin toss leads one to believe that $2+2=5$ rather than that $2+2=4$). We thus get a way of accounting for how a belief even in a necessary truth might be veritically epistemically lucky, and hence why safety, construed through the lens of anti-luck epistemology, needs to be understood in a particular way.

II. THE MODAL ACCOUNT OF RISK

Even despite the attractions of anti-luck epistemology as just set out, however, I nonetheless want to urge that we need to modify the view. In order to see why, we need to consider the notion of risk, which as we will see is closely related to luck.

There is clearly a close relationship between the notions of risk and luck. Suppose, for example, that one opts to board a rickety-looking plane, and after an eventful and clearly dangerous flight, emerges unscathed. We would naturally judge that one is lucky to survive this flight and also that taking the flight was high-risk. Indeed, it seems that these assessments only really differ in terms of whether they are forwards-looking or backwards-looking.¹⁴ That is, risk assessments seem to be essentially forwards-looking: if you take that flight, then you are at risk. Luck assessments, in contrast, seem to be essentially backwards-looking: you were lucky to have survived that flight. Nonetheless, what gives rise to the luck and the risk in both cases is the same—*wiz*, the modal closeness of the scenario in which this plane crashes and kills you. Thus, luck and risk seem to be very closely related notions, differing only in terms of the direction of assessment.

The empirical literature backs this up, in that there is a lot of empirical support for the idea that subjects' judgements about risk and luck tend to go hand-in-hand. In particular, just like luck, subjects' judgements about risk track the modal closeness of the target event rather than its probabilistic likelihood. So subjects might grant that the probabilistic likelihood of two events is broadly the same, and yet nonetheless characterize one of them as being riskier than the other because they regard this event as modally closer.

¹⁴ I am grateful to Jesus Navarro who first pointed this out to me.

A good example of this is subjects' judgements about the risks involved in various kinds of transport. While subjects will grant that the probability of sustaining serious injury when, say, driving a car is much, much higher than alternative forms of transport, such as taking the train, they nonetheless tend to judge that car driving is not an especially risky activity (i.e., no more risky, or at least not especially riskier, than taking the train). There are various explanations for this. It is certainly true, for example, that cognitive biases have a role to play in leading subjects to make these assessments of risk. The fact that one is driving one's car, as opposed to being a passenger (as on a train), makes the 'illusion of control' bias relevant, for example.¹⁵ This leads subjects to overestimate their control over events associated with car driving, such as their propensity to have accidents. This bias, coupled with the fact that subjects tend to overestimate their expertise (most people think that they are above-average drivers),¹⁶ leads them to regard driving a car as a not especially risky activity, even taking the relatively high probability of car accidents (when compared with some other forms of transport) into account.¹⁷

We can explain what is going on here in terms of our account of luck. The car driver in the grip of these cognitive biases has a conception of the actual world such that the possible world in which they incur serious injury while driving is not especially close, and hence not a serious risk (even despite the relatively high probabilities in play). That is, keeping fixed salient initial conditions (the subject's above average driving skill, for example), it is not a matter of luck that they avoid serious accident on a given car journey. Judgements about luck thus dovetail with judgements about risk. To say that a target event is risky is to say that (keeping relevant initial conditions for that event fixed) it obtains in close possible worlds. As the modal distance between the actual world and the possible world where the target event obtains becomes more remote, so

¹⁵ Indeed, subjects tend to judge travelling by car as riskier when it is made clear that they will be a passenger rather than the driver. See, for example, F. P. McKenna, 'It Won't Happen To Me: Unrealistic Optimism or Illusion of Control?', *British Journal of Psychology* 84 (1993), 39-50. For more on the illusion of control, see E. J. Langer, 'The Illusion of Control', *Journal of Personality and Social Psychology* 32 (1975), 311-28; S. C. Thompson, 'Illusions of Control: How We Overestimate Our Personal Influence', *Current Directions in Psychological Science* 8 (1999), 187-190; and Thompson, 'Illusions of Control', *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*, (ed.) R. F. Pohl, 115-25, (Hove, UK: Psychology Press, 2004).

¹⁶ This is the so-called *overconfidence bias*. Famously, in a US study it was found that 93% of drivers rated their driving abilities as above average. See O. Svenson, 'Are We Less Risky and More Skillful Than Our Fellow Drivers?', *Acta Psychologica* 47 (1981), 143-51. Indeed, interestingly (though perhaps not surprisingly), those with low levels of skill are often *more* apt to overestimate their skill levels, a phenomenon known as the 'Dunning-Kruger effect'. See J. Kruger and D. Dunning, 'Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments', *Journal of Personality and Social Psychology* 77 (1999), 1121-34.

¹⁷ Note that it doesn't matter for our purposes that we are dealing here with risk ascriptions that are the product of cognitive bias (any more than it mattered when we appealed to similar empirical literature in support of the modal account of luck). This is because the point in play concerns what agents are *trying* to track in their risk judgements, and not whether they are successful in this regard. I expand on this point in some detail in Pritchard, 'The Modal Account of Luck', *op. cit.*, §2. For some useful empirical studies on risk perception, see P. Slovic, 'Perception of Risk', *Science* 236 (1987), 280-85; L. Sjöberg, 'Factors in Risk Perception', *Risk Analysis* 20 (2000), 1-11; and Sjöberg, B.-E. Moen and T. Rundmo, *Explaining Risk Perception: An Evaluation of the Psychometric Paradigm in Risk Perception Research*, (Trondheim, Norway: Rotunde, 2004).

the riskiness of the event lessens. At some point, the target event is so modally remote as to not be significantly risky, and hence we tend to judge that there is no risk involved.¹⁸

That said, although it is possible to accommodate these judgements about risk in terms of the modal account of luck, we lose some theoretical purchase on the topic in hand if we do so. There are a number of reasons for this. First, notice that in accounting for our judgements about risk in terms of the modal account of luck, we are implicitly applying the modal account of luck in a very specific way. When we make judgements about the extent of risk involved in an event (or activity etc.), our focus is not merely on whether that event obtains in the modal neighbourhood as it is with luck, but rather on a specific unwelcome event associated with that event, what we can call the *risk event*. So, for instance, when taking a flight on a plane and assessing the risk involved, our specific interest is on the modal closeness of the relevant risk event—in this case, for example, this will likely be the scenario where the plane crashes.

This leads us to a second salient difference between our judgements about luck and about risk, in that the latter specifically concerns ‘negative’ events, such as hazards (e.g., the risk of being in a car accident), while we can quite comfortably think of the former in terms of positive events like lottery wins. So, while it is a matter of luck that one wins a lottery, this does not normally have anything to do with assessments of risk because lottery wins are welcome events. In contrast, imagine a lottery draw where the outcome is, unusually, unwelcome. Perhaps, for example, one has the misfortune to be in a ‘Hunger Games’ type scenario where the ‘winner’ of the lottery draw will have to endure some terrible fate. Now it would make sense to use the language of risk. After all, there is now a close possible world where the target risk event—i.e., that one’s numbers are drawn—obtains.

There is another reason for keeping our accounts of luck and risk apart, and this is that where these proposals coincide in extension, our interest in excluding luck is usually to be explained by our interest in excluding risk, rather than *vice versa*. That is, risk is arguably the more fundamental notion in play here. Think again about the ‘Hunger Games’ style scenario just described. Why would we care about the fact that it is just a matter of luck that one could be chosen for this terrible fate, such that being chosen is an easy possibility? Well, the natural explanation to offer is that such luckiness is a concern precisely because it indicates that there is a high risk that this unwelcome event might befall one. That is, it is the concern with avoiding high risk that is explaining the concern with luck in such a case, rather than *vice versa*.

There are thus advantages to keeping the modal account of luck and the modal account of

¹⁸ We are then in the region of what is known in the literature as *de minimis* risks. For discussion, see J. Mumpower, ‘An Analysis of the *De Minimis* Strategy for Risk Management’, *Risk Analysis* 6 (1986), 437–46, and M. Petersen, ‘What Is a *De Minimis* Risk?’, *Risk Management* 4 (2002), 47–55.

risk apart, despite their superficial similarities, and despite the fact that in a wide range of cases ascriptions of luck and ascriptions of risk will tend to go hand-in-hand. According to the *modal account of risk*, the level of risk involved is determined by how modally close the target risk event is. In particular, where the risk event is modally close, then it is high-risk (even if it is a probabilistically unlikely event), but where it is not modally close, then it is low-risk. As with the modal account of luck, in making these assessments we need to keep relevant initial conditions fixed across cases. So, when evaluating the risk involved in catching a plane (i.e., the risk of a plane crash), we are interested in those possible worlds where one continues to catch the plane, since it clearly wouldn't be relevant to this particular assessment of risk that in those close possible worlds where one doesn't catch the plane one isn't involved in a plane crash. Moreover, as with the modal account of luck, we have a continuum picture of risk in play here, with an intolerance to high levels of risk (i.e., where the risk event is modally close), a tolerance to low levels of risk (i.e., where the risk event is modally distant), and a continuum of increasing tolerance in between.¹⁹

III. ANTI-RISK EPISTEMOLOGY

With this distinction between the modal account of luck and the modal account of risk in play, we can now evaluate the prospects for an anti-risk epistemology, and in particular the prospects for such a view in contrast to an anti-luck epistemology. The first thing to note is that the differences between the two views will not be radical since, as we have noted above, there is a considerable overlap in our judgements about luck and risk, and hence there is also a considerable degree of common ground when it comes to the modal accounts of luck and risk. Accordingly, we would expect epistemic luck and epistemic risk to go hand-in-hand too. Generally speaking, a belief that is subject to veritic epistemic luck will have been formed in an epistemically risky fashion, and *vice versa*.

Nonetheless, there are important differences between an anti-luck epistemology and an anti-risk epistemology. In order to get a handle on anti-risk epistemology, and how it differs from anti-luck epistemology, it will be useful to begin by noting a theoretical lacuna in anti-luck epistemology that is usually glossed over. Recall that the luckiness of an event depends on how

¹⁹ For a more detailed defence of the modal account of risk, see Pritchard, 'Risk', *Metaphilosophy* 46 (2015), 436-61. For an explanation of how the modal account of risk can be applied to two prominent debates in legal philosophy, see Pritchard, 'Legal Risk and the Arithmetic of Criminal Justice', *Jurisprudence* (forthcoming). For two useful overviews of the philosophy of risk, see S. O. Hansson, 'Philosophical Perspectives on Risk', *Techné: Research in Philosophy and Technology* 8 (2004), 10-35, and 'Risk', *Stanford Encyclopedia of Philosophy*, (ed.) E. N. Zalta, (2014) [available at: <http://plato.stanford.edu/archives/spr2014/entries/risk/>]. Note, however, that Hansson's focus in these articles (in common with most contemporary treatments of risk) is on a probabilistic conception of risk, in contrast to the modal account defended here.

modally close the non-obtaining of that event is. When the modal account of luck is applied to veritic epistemic luck, this gets translated in terms of a true belief (= cognitive success) where it is an easy possibility that one forms a false belief (= cognitive failure). If we were to follow the modal account in a strict fashion, however, then veritic epistemic luck should be a true belief where it is an easy possibility that one doesn't form a true belief. This is a weaker claim, in that not forming a true belief does not entail forming a false belief. In short, the notion of 'cognitive failure' in play here is ambiguous, in that it can either mean the absence of cognitive success (of which false belief is one variant), or it can specifically mean false belief. The question is thus why anti-luck epistemology opts for the more restrictive rendering of the anti-luck condition (i.e., where cognitive success is translated as false belief, as opposed to the disjunction of false belief and non-belief), given that this restrictive rendering is not in fact motivated by the modal account of luck?

The reason why becomes apparent once one follows through the implications of this move, in that it would rule out genuine cases of knowledge. Imagine that an agent just happens to be in the right place and at the right time to overhear a crucial snippet of conversation, and thereby forms a true belief on this basis. Given that the opportunity presented to the agent was properly exploited, this true belief ought to amount to knowledge. Even so it could well be that in all close possible worlds this agent wouldn't be so situated as to form this true belief, in that she would instead not form *any* belief on this basis since she was never well enough situated to hear enough of the conversation to work out what was being said. In this case while it is not an easy possibility that she forms a false belief, it is an easy possibility that she forms no belief. Accordingly, if one opts for the weak inclusive rendering of the anti-luck condition (i.e., where cognitive failure means either non-belief or false belief), then one would be obliged to treat the agent as lacking knowledge.

The inclusive rendering of the anti-luck condition thus seems to be generating entirely the wrong result here. While it is undoubtedly lucky that the agent ended up with a true belief, this is not the kind of epistemic luck that is incompatible with knowledge. It is rather epistemic luck of a purely evidential variety, in that the agent was lucky to have the evidence that she did. Given that she has this evidence, however, she is in a position to exploit it to gain knowledge. If we insist that epistemic luck of this kind is incompatible with knowledge, then we will be obliged to rule out a lot of cases of this type as being genuine instances of knowledge, contrary to intuition. (Indeed, we will be obliged to claim that there can be no such thing as a lucky discovery).²⁰

It is thus very important that anti-luck epistemology opts for the restrictive rendering of the

²⁰ For more on the notion of evidential epistemic luck, particularly *qua* a benign form of epistemic luck which (unlike veritic epistemic luck) is compatible with knowledge, see Pritchard, *Epistemic Luck*, *op. cit.*, ch. 5.

anti-luck condition, rather than the inclusive rendering that is generated by appealing to the modal account of luck. This theoretical lacuna in anti-luck epistemology is not immediately apparent because it generates a basis-relative version of safety (i.e., one where we keep the original basis for belief fixed across possible worlds). As noted above, there is a rationale for this restriction, in that the modal account of luck demands that we keep the relevant initial conditions for the target event fixed, and keeping the basis for belief fixed is simply the cognitive analogue of this restriction. This disguises the problem in play, however, since in the normal cases of evidential epistemic luck, such as the one just given, the subject's basis for belief is also changing, and hence they do not arise as direct counterexamples to anti-luck epistemology. Accordingly, once basis-relativity is taken into account it might not seem to matter whether one opts for the restrictive or the inclusive rendering of the anti-luck condition (on this score anyway), since one cannot generate a direct counterexample to the view either way.

Nonetheless, once we have a handle on the structure of the problem in play here, it ought to be clear how one can adapt the cases to resurrect a counterexample to anti-luck epistemology on the inclusive reading of the anti-luck condition (and hence force a wedge between the two formulations of the view). Imagine, for example, a subject who has a sound memorial basis for her 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 she is disposed to doubt herself. Let's say it is just a psychological fact about this subject that her confidence in what she believes is quite variable, even though there is no epistemic basis for this variability (it is not as if, for example, she has reasons to doubt her memory). It will now be true of such a subject that in close possible worlds she will have the same basis for belief as in the actual world and yet will not believe the target proposition on this basis. Accordingly, by the lights of the inclusive reading of anti-luck epistemology (but not the restrictive reading) this agent ought to be regarded as having a true belief which is subject to veritic epistemic luck and hence not in the market for knowledge. But this seems like just the wrong result. It is not, after all, as if the subject is in error in close possible worlds. In particular, in the relevant sense of cognitive failure that interests us—i.e., false belief, formed on the same basis as in the actual world—the subject has exhibited no cognitive failure at all, and hence ought to have a belief that is in the market for knowledge (at least as far as the anti-luck platitude goes anyway).

This is where the shift away from anti-luck epistemology to anti-risk epistemology becomes important. Recall that the modal account of risk is not concerned merely with the non-obtaining of the target event in the modal neighbourhood, but rather more specifically with whether the relevant risk event obtains in the modal neighbourhood. In terms of eliminating epistemic risk, this will translate as the claim that one is looking for a true belief on a particular basis (= cognitive

success) where the risk event of forming a false belief on that same basis (= cognitive failure) is not too modally close. In short, we are demanding that knowledge excludes *veritic epistemic risk*, which is the analogue of veritic epistemic luck in an anti-luck epistemology. By appealing directly to anti-risk epistemology we can thus close this hitherto unnoticed lacuna in anti-luck epistemology by appealing directly to the risk event of false belief.

As with anti-luck epistemology, the guiding thought will be that we are intolerant of epistemic risk which is very high (i.e., where the risk event of false belief is modally very close), and hence will not ascribe knowledge, but we will be tolerant of epistemic risk that is very low (i.e., where the risk event of false belief is modally far-off). In between these two extremes, there will be a continuum of increasing tolerance of epistemic risk as the risk event moves further modally afield.

A second advantage of anti-risk epistemology over anti-luck epistemology relates to the fact that, as noted above, risk seems a more basic notion than luck, in that we naturally explain a concern to eliminate the latter in terms of a concern to eliminate the former rather than *vice versa*. This is particularly relevant to anti-luck epistemology because of how the anti-luck platitude has been questioned in some quarters in the recent literature. Indeed, some have even gone so far as to argue that there is no essential tension between knowledge and veritic epistemic luck.²¹ If one expresses the incompatibility in question purely in terms of luck, then one might see some mileage in this claim on account of the fact that it at least looks like an open question whether knowledge could be compatible with (veritic epistemic) luck. But if we now understand that concern to eliminate (veritic epistemic) luck in terms of a desire to eliminate high levels of epistemic risk, then the question doesn't look nearly so open. Risk events, after all, are by their nature unwelcome events, so it is hardly surprising that we desire to eliminate high levels of risk. Insofar as the risk event of false belief is what is associated with knowledge, then it follows that we should not be surprised that we have a strong intuition that knowledge ought not to be compatible with high levels of epistemic risk. In short, the claim that knowledge is incompatible with veritic epistemic risk should strike us as far more compelling than the claim that knowledge is incompatible with veritic epistemic luck. Once we see that denying the latter means also denying the former, then that ought to remove much of the temptation for maintaining that knowledge can co-exist with veritic epistemic luck.

Now one might think that the import of this point to the contemporary debate in epistemology is fairly narrow, given that very few epistemologists argue that knowledge and veritic

²¹ For a very clear example of this kind of thesis, see Stephen Hetherington, 'There Can be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Edn.), (eds.) M. Steup & J. Turri, 164-76, (Oxford: Blackwell, 2013).

epistemic luck are entirely compatible. Interestingly, however, while few epistemologists are willing to concede that knowledge is compatible with veritic epistemic luck in general, there is a growing band of epistemologists who argue that knowledge is compatible with certain kinds of veritic epistemic luck. If the foregoing is correct, however, then even this more modest claim entails embracing the compatibility of knowledge with high levels of epistemic risk, and hence ought to be similarly problematic.

The issue that concerns us here relates to a distinction that I have elsewhere drawn between two kinds of veritic epistemic luck: *intervening epistemic luck* and *environmental epistemic luck*.²² The first kind is more usual, and is the type of veritic epistemic luck that one typically finds in standard Gettier-style cases. So, for example, imagine a farmer in a field, in good cognitive conditions, looking at what appears to be a sheep, and so forming on this basis the belief that there is a sheep in the field. Suppose, however, that what the farmer is in fact looking at is not a sheep at all but rather a big hairy dog. Nonetheless, her belief is true even so, since behind the big hairy dog, hidden from view, is a sheep.²³ We thus have a true belief which is skillfully formed (and hence which is justified), but which doesn't amount to knowledge because of the veritic epistemic luck involved. In particular, given how the belief was formed, it could so very easily have been a false belief (i.e., if the sheep hidden from view had wandered out of the field, but everything else had stayed the same).

In the case just described, something in effect *intervenes* between the subject's belief and the target fact. It is not a genuine sheep that the subject is looking at, but something else entirely, even though her belief is true regardless. This is thus an example of veritic epistemic luck that specifically involves intervening epistemic luck. Significantly, however, not all cases of veritic epistemic luck involve intervening epistemic luck.

Consider instead the famous barn façade case.²⁴ Here we have a subject who encounters a genuine barn in conditions that are apt for seeing a barn. Crucial to the example, however, is that the subject is in barn façade county where just about everything else that looks like a barn is not a genuine barn at all but rather a fake. In this case, nothing intervenes between the agent's belief and the target fact. Even so, the belief in question is subject to veritic epistemic luck, since it is a true

²² I first drew this distinction between environmental and intervening epistemic luck in Pritchard, *Knowledge*, (London: Palgrave Macmillan, 2009), chs. 3-4. For further discussion of this distinction, see Pritchard, 'Knowledge, Understanding and Epistemic Value', *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O'Hear, 19-43, (Cambridge: Cambridge University Press, 2009), and 'Anti-Luck Virtue Epistemology', *op. cit.*; and Pritchard, A. Millar and A. Haddock, *The Nature and Value of Knowledge: Three Investigations*, (Oxford: Oxford University Press, 2010), chs. 2-4. See also Jesper Kallestrup and Pritchard, 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy* 22 (2014), 335-57.

²³ This case was originally due to Roderick Chisholm, *Theory of Knowledge*, (2nd ed.), (Englewood Cliffs, NJ: Prentice-Hall, 1977), 105.

²⁴ See Alvin Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73 (1976), 771-91, who credits the example to Carl Ginet.

belief that could so very easily have been a false belief, given how it was formed (i.e., had the subject happened to have been looking at one of the fake barns nearby instead). The source of the veritic epistemic luck is thus entirely in the environment, as opposed to intervening between belief and fact, and hence this is a case of veritic epistemic luck that specifically involves environmental epistemic luck.

This distinction between two types of veritic epistemic luck is significant for our purposes because while very few epistemologists are willing to grant that knowledge is compatible with veritic epistemic luck in general, there is a growing constituency of epistemologists who are at least sanguine about the idea that knowledge might be compatible with environmental epistemic luck.²⁵ One can see the temptation to take this route. After all, there is plausibly a greater degree of epistemic support in favour of the target belief in cases of environmental epistemic luck when compared with intervening epistemic luck. For example, cases of environmental epistemic luck arguably involve a cognitive success that is creditable to the subject's cognitive agency, something which does not apply in cases of intervening epistemic luck. That is, I don't think anyone is tempted to say that the farmer's cognitive success in the example just offered is creditable to her cognitive agency. It is, instead, attributable to the happenstance there is a sheep in the field behind the sheep-shaped object that she is looking at. In contrast, there is at least a case to be made that the subject in the barn façade case has a true belief that is creditable to her cognitive agency (i.e., to her exercise of relevant barn-spotting abilities).²⁶ If that's right, then there is a positive epistemic standing available to beliefs that are subject to environmental epistemic luck that is not available to corresponding beliefs that are subject to intervening epistemic luck. Not all forms of veritic epistemic luck are thus on an epistemic par.

Even granted this point, however, I think we should resist the urge to attribute knowledge in cases of environmental epistemic luck. Note first that the positive epistemic standing that may be available in cases of environmental epistemic luck is really by-the-by. One can, after all, imagine two beliefs subject to intervening epistemic luck where one of the beliefs is more justified than the other—this epistemic difference wouldn't incline us to regard the more justified belief as a potential instance of knowledge. The key question is thus whether the relevant barrier to knowledge that is in play when we engage with the problem of veritic epistemic luck has been cleared.

²⁵ See Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, (Oxford: Oxford University Press, 2007), ch. 5, for a particularly high-profile endorsement of this kind of line (though note that Sosa's focus is not on the barn façade case here, but rather the structurally similar 'jokester' case). I critically discuss Sosa's reasons for ascribing knowledge in this case in Pritchard, 'Apt Performance and Epistemic Value', *Philosophical Studies* 143 (2009), 407-16.

²⁶ This is especially so if the barn façade case is modalised—i.e., such that there is in fact no deception taking place in the actual world, but only in close possible worlds. For further discussion of this point, see Pritchard, 'Anti-Luck Epistemology and the Gettier Problem', *op. cit.* See also the *epistemic twin earth case* as described by Kallestrup and Pritchard, *op. cit.*

I think that casting this question specifically in the language of veritic epistemic risk helps us to see that it hasn't. After all, even despite this putative difference in epistemic standing, it remains the case that beliefs that are subject to environmental epistemic luck are no less subject to veritic epistemic risk than their counterparts which are subject to intervening epistemic luck. In particular, in both cases the risk event of false belief formed on the same basis is modally very close. This exposes the radical nature of what is being proposed by those who allow knowledge to be compatible with environmental epistemic luck. So expressed, the stance can appear less radical than it is. It is only when we re-cast this point in terms of veritic epistemic risk—i.e., as the claim that knowledge can be compatible with high levels of epistemic risk—that we recognise just how radical this claim is. The appeal to a particular kind of veritic epistemic luck is thus a red herring. While claiming that only some kinds of veritic epistemic luck are compatible with knowledge can seem like a much more modest thesis than the claim that they all are, the fact remains that the putative compatibility of knowledge and environmental epistemic luck entails that knowledge is compatible with high levels of epistemic risk. So construed, there is nothing at all modest about such a proposal.

Given the foregoing, one might wonder why epistemologists have focused on an anti-luck platitude with regard to knowledge rather than going directly for an anti-risk platitude. The short answer to this question is that I'm not sure, but I do have a hunch. I think it relates to the point I noted earlier about the different directions of assessment involved when we evaluate something in terms of risk and luck. The big debate about epistemic luck in epistemology arose in response to the Gettier problem, which notoriously concerns a justified yet luckily true belief. But since the belief is also formed in an epistemically risky fashion, then why the focus on luck here rather than risk? One possible explanation is that the formulation of the Gettier problem invites us to take a backwards-looking perspective on the case, on account of the fact that that it is stipulated in advance that the belief in question is true. Just as once one has stepped off the rickety plane that almost killed you, one naturally tends to focus on the good luck involved in surviving the journey rather than the risk in play throughout, the same seems true here. Given that things have turned out well, our natural focus is on the luck involved rather than the risk of error. But this shouldn't blind us to the fact that both are present, and that our real concern with the former is due to how it indicates the presence of the latter.

In any case, the upshot of the foregoing is that the differences between anti-luck epistemology and anti-risk epistemology, while small, are nonetheless theoretically significant. Those inclined towards the former project should thus realign themselves with the new rendering of this project offered by the modal account of risk. And those who are opposed to anti-luck epistemology would be wise to reflect on how their objections to this view fare once it is

reconfigured along anti-risk lines.²⁷

DUNCAN PRITCHARD

University of Edinburgh

²⁷ I am grateful to anonymous reviewers for the *Journal of Philosophy* for comments on an earlier version of this paper. Special thanks to Jesus Navarro for long discussions about this article, and thanks also to Modesto Gomez Alonso, Zenon Bankowski, Laura Candiotta, J. Adam Carter, Alkis Kotsonis, Lex Lai, Tony Lau, Lukas Schwengerer, Kegan Shaw, Jacques Vollet, and Di Yang. Some of this material was presented at conferences at the University of Leeds and the University of Edinburgh during 2015, and at a postgraduate reading group at the University of Edinburgh during 2016, and I am grateful to the audiences for their comments.