

Anti-risk epistemology and negative epistemic dependence

Duncan Pritchard¹ 

Received: 6 March 2017 / Accepted: 3 October 2017 / Published online: 20 October 2017
© The Author(s) 2017. This article is an open access publication

Abstract Support is canvassed for a new approach to epistemology called *anti-risk epistemology*. It is argued that this proposal is rooted in the motivations for an existing account, known as *anti-luck epistemology*, but is superior on a number of fronts. In particular, anti-risk epistemology is better placed than anti-luck epistemology to supply the motivation for certain theoretical moves with regard to safety-based approaches to knowledge. Moreover, anti-risk epistemology is more easily extendable to epistemological questions beyond that in play in the theory of knowledge specifically. A key advantage of the view, however, is that anti-risk epistemology fares much better than anti-luck epistemology when it comes to accounting for the phenomenon of *negative epistemic dependence*. In particular, anti-risk epistemology is ideally placed to explain why such epistemic dependence is incompatible with knowledge, even when the negative epistemic dependence in play is of a purely modal variety (as one finds in *epistemic twin earth cases*).

Keywords Epistemology · Epistemic luck · Epistemic risk · Epistemic dependence · Knowledge · Epistemic twin earth cases

1 Anti-luck epistemology

One of the guiding platitudes in epistemology is that knowledge excludes luck.¹ It is an intuition of this general sort that is meant to explain, for example, why knowledge

¹ For a dissenting voice in this regard, see [Hetherington \(2013\)](#), which is in turn a response to [Pritchard \(2013\)](#).

✉ Duncan Pritchard
duncan.pritchard@ed.ac.uk

¹ University of Edinburgh, Edinburgh, UK

is lacking in Gettier-style cases, even though the belief is true and justified (and often enjoys other forms of epistemic support too, such as being formed via a reliable basis, etc). The guiding thought behind an *anti-luck epistemology* is that such a platitude presents one with a principled way of approaching the theory of knowledge. Rather than working piecemeal through a range of putative anti-Gettier conditions and seeing which one can deal with the cases in play, we instead adopt a three-stage approach. The first stage is to offer an account of luck. The second stage is to specify the particular way in which knowledge excludes luck (there can be lucky discoveries, after all, so there must be ways in which there can be lucky knowledge). The third and final stage is to put these two parts together and thereby come up with a formulation of an anti-luck condition on knowledge.

In terms of the first stage, I have argued at length elsewhere for a *modal account of luck*.² Very roughly, on this view lucky events are to be understood as those events that obtain in the actual world but which fail to obtain in close possible worlds where the relevant initial conditions for the target event are kept fixed.³ Winning the lottery is thus lucky, since in close possible worlds where the relevant initial conditions for the event are kept fixed (e.g., the lottery is still fair and has similarly long odds), one fails to win. In contrast, it is not a matter of luck that the parachute failed to open given that it was sabotaged by the enemy spy, expert at such activities. In all close possible worlds where we keep the relevant initial conditions for this event fixed (i.e., the enemy spy continues to be highly skilled in this regard, and continues to sabotage the parachute), the parachute similarly fails to open.

Moreover, we can delineate the luckiness of an event in terms of how modally close its non-obtaining would be, such that we can determine a continuum of degrees of luck. So, for example, suppose that one is nearly hit by a sniper's bullet, which whizzes by one's ear, such that one is lucky not to be shot. *Ceteris paribus*, this is luckier than an equivalent scenario where the bullet flies by a metre or so away, which is in turn luckier than an equivalent scenario where the bullet flies by some further distance away, and so on. At some point, as the modal vicinity of one's being hit by the bullet becomes ceases to be all that close, we would cease to regard the event in question as being lucky at all.⁴

In terms of the second stage, it is widely accepted that the kind of epistemic luck that is incompatible with knowledge is what is known as *veritic epistemic luck*. This occurs when it is a matter of luck that one's belief is true, given how one formed it.

² See especially Pritchard (2014). See also Pritchard (2004, 2005, 2006). Although this is the dominant theory of luck in the philosophical literature, it does have its detractors, and there are proponents of alternative views. For a useful recent volume which collects a range of literature on the philosophy of luck, including pieces that are both in favour of, and critical of, the modal account, see Pritchard and Whittington (2015). See also Levy (2011) and Coffman (2015).

³ I am here characterizing possible worlds in the standard way in terms of a similarity ordering. See, for example, Lewis (1973, 1987). For further defence of this account of possible worlds in the context of the modal account of luck, see Pritchard (2014, Sect. 3).

⁴ There is a wealth of empirical work with regard to the psychology of luck ascriptions that supports the modal account. See Kahneman and Varey (1990), Teigen (1995, 1996, 1997, 1998a, b, 2003), Tetlock (1998), and Tetlock and Lebow (2001). See Pritchard and Smith (2004) for a survey of this research. For also Pritchard (2014), (2016b, §1).

Veritic epistemic luck is contrasted with epistemically benign forms of epistemic luck that are compatible with knowledge, such as the *evidential epistemic luck* that one happens to be in the right place and time to have the excellent evidence that one does. Provided that one properly exploits this epistemic advantage, then one can come to have knowledge regardless (i.e., in a way that is not, in particular, subject to veritic epistemic luck).⁵

We now come to stage three, which is putting the modal account of luck together with an account of knowledge such that it is incompatible with veritic epistemic luck. The result is a modal condition on knowledge. In particular, when one knows one's true belief is formed in such a way that one couldn't have easily on that same basis formed a false belief. Moreover, in line with our continuum conception of degrees of luck, we will tend to be completely intolerant of veritic epistemic luck when it involves forming a false belief on the very same basis in very close possible worlds, with that tolerance increasing as we move further out, to the point where we will eventually be sanguine about the possibility of error because it is so remote.

Interestingly, the kind of modal condition that anti-luck epistemology is here motivating is essentially a version of the *safety* condition for knowledge.⁶ This is usually informally understood as the claim that one's true belief could not have easily been false. We are now able to offer a way of fleshing out this notion in modal terms, where we bring in basis-relativity (in order to keep the initial conditions for the target event, in this case the true belief, fixed), and where we also effectively 'weight' the very closest possible worlds over not-so-close possible worlds that are also in the local neighbourhood, in order to accommodate the continuum aspect of luck. It is thus not just that anti-luck epistemology generates independent support for a particular condition on knowledge, but also that it gives us a principled basis for unpacking how best to understand the usual informal presentation of that condition.⁷

Anti-luck epistemology also fares well in terms of accounting for the kinds of cases that epistemologists have claimed involved knowledge-undermining epistemic luck. Gettier-style cases, after all, essentially involve a subject forming a belief in a fashion which, although it might enjoy certain kinds of epistemic pedigree (e.g., it is justified, or reliably formed, say), is nonetheless such that, given how it was formed, it is a close possibility that the subject ends up with a false belief.

Or consider the lottery puzzle, which also trades on knowledge-undermining epistemic luck. Why is it that it seems one can know that one has lost the lottery by checking the results in a reliable newspaper (even when one is oblivious as to the long odds involved in winning), but one cannot know that one has lost the lottery by merely

⁵ For further discussion of epistemically malignant and epistemically benign varieties of epistemic luck, see Pritchard (2004, 2005).

⁶ For some key defences of the safety principle, see Luper (1984, cf. 2003), Sainsbury (1997), Sosa (1999), Williamson (2000), and Pritchard (2002).

⁷ In particular, anti-luck epistemology favours the safety condition on knowledge over the other main competing anti-luck condition in the literature, which is the *sensitivity* condition, at least as the latter is usually understood anyway. For more on this point, see Pritchard (2007, 2012b). For some key defences of the sensitivity principle, see Dretske (1970, 1971), Nozick (1981), Roush (2005), Becker (2007), Black and Murphy (2007), and Black (2002, 2008). For a comparative overview of the safety and sensitivity principles, see Pritchard (2008) and Black (2011).

reflecting on the long odds involved in winning? After all, from a purely probabilistic point of view, it seems one ought to be better placed to know in the latter case, as the odds of success are massively in one's favour (even more so than in the former case).

Anti-luck epistemology offers a very straightforward explanation of what is going on here, which is that our judgements about knowledge and luck are ultimately sensitive less to the probabilities involved and more to the modal closeness of error. In particular, it is important to remember that low probability events can nonetheless be modally close events, in that very little need change about the actual world in order for them to occur. A lottery win is a case in point, in that although this is very low probability, precisely the reason why people play lotteries is that they could very easily occur—just a few coloured balls need to fall in a different configuration.⁸ This is why forming one's true belief that one has lost the lottery merely by reflecting on the odds involved is not knowledge, since there is a very close possible world where one forms one's belief on the same basis and believes falsely as a result. In contrast, given the lengths reliable newspapers go to ensure that they print the right lottery results, the possible world in which one forms a false belief on this basis is not very close at all, which is why this can be a route to knowledge.

The foregoing is not meant to be a full defence of anti-luck epistemology, since that would take us too far afield for our purposes, but merely an overview of its virtues.⁹ In fact, I want to argue that while anti-luck epistemology is broadly on the right lines, it nonetheless needs to be reconfigured in an important respect. In particular, we need to alter it so that the notion of *epistemic risk* is central, rather than epistemic luck. As we will see, this reconfiguration offers several main advantages: it removes a couple of theoretical lacunae in anti-luck epistemology, it offers us a better diagnostic account of why knowledge has such a distinctive modal profile, and the resulting proposal is more easily extendable to other epistemic standings beyond knowledge. These considerations will be set out in the next section. But there is also another key advantage, which will be our focus in Sect. 3, which how an anti-risk epistemology is much better placed to account for why cases of what I refer to as *negative epistemic dependence* are knowledge-undermining.

⁸ That is, placing a bet on a lottery win (i.e., by purchasing a ticket) is very unlike the usual case of placing a bet on an event with astronomically low odds (e.g., that a completely useless football team will win this year's FA Cup). Interestingly, the slogan for the UK's national lottery "It could be you!" confirms this point. This is clearly not the 'could' of probability, since in this sense it (realistically) couldn't be you, as the odds are astronomically against 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 (such as the completely useless team winning the FA Cup) would be, from a rational point of view, much worse.

⁹ I offer a fuller defence of anti-luck epistemology in a number of places see especially Pritchard (2007, 2012b).

2 From anti-luck epistemology to anti-risk epistemology

In order to understand why the shift from an anti-luck to an anti-risk epistemology might be beneficial, we first need to notice some of the key differences between luck and risk. As I've argued elsewhere, luck and risk are closely entwined notions, and as such risk is just as amenable to a modal account, cast along broadly the same lines as the modal account of luck.¹⁰ Indeed, in the relevant empirical literature, work on luck ascriptions and risk ascriptions often goes hand-in-hand.¹¹ This means that just as we are to think of luck in modal terms by focussing on relevant features of the modal environment, so we are to think of risk in modal terms by focussing on essentially similar features of the modal environment.¹² With these points in mind, it is easier to specify when these notions diverge and thus where the modal account of risk diverges, in detail, from the modal account of luck rather than where they converge. I will here focus on two key differences.¹³

The first is that assessments of risk involve an essentially forwards-looking perspective on events, unlike luck assessments, which involve an essentially backwards-looking perspective.¹⁴ Take our agent from earlier who narrowly missed being hit by the sniper's bullet. In terms of a risk assessment of this scenario, we take the perspective of the subject *prior* to the bullet flying by and observe that she was at a high risk of being hit. In terms of a luck assessment of this scenario, in contrast, we take the perspective of the subject *after* the bullet has missed and observe that she was lucky not to have been hit.

A related point in this regard is that in cases where we are concerned with an assessment pairing of luck and risk, such as in the sniper's case just considered, I take it that normally the reason why we are interested in luck is because of the risk involved,

¹⁰ See especially Pritchard (2015b).

¹¹ For more on this point, see Pritchard (2004, 2015b); §4). See also the empirical literature cited in footnote 4 above, much of which simultaneously concerns both luck and risk ascriptions.

¹² Note that the modal account of risk runs counter to the most prominent theory of risk in the literature, which is cast along broadly probabilistic lines. See Hansson (2004, 2014) for two excellent surveys of recent theoretical work on risk (which puts the probabilistic account centre-stage). See (Pritchard 2015b, §§2–3) for a comparative discussion of the modal and probabilistic accounts of risk.

¹³ Note that these two differences are far from exhaustive. For example, one further difference is that luck is more clearly a threshold notion. In terms of the modal account of luck, this means that once the possible world where the non-obtaining of the target event becomes sufficiently far-out, then it no longer makes sense to talk of that event as being lucky at all. In contrast, we can usually meaningfully talk of a very low level of risk, even if that risk is so negligible that we wouldn't want to talk about the target event being risky [though there is usually a minimal threshold for risk what is known in the literature as *de minimis risk*. See Mumpower (1986) and Petersen (2002) for discussion]. Another difference is that luck is *bivalent*, in the sense that there can be both good and bad luck. In contrast, risk is often (though not exclusively) negatively characterised in terms of the avoidance of hazardous events. See Pritchard (forthcoming) for a discussion of those unusual cases where risk adds positive value to an activity, such as dangerous sports or when an artist takes aesthetic risks. See Pritchard (2015b) for further discussion of the differences between luck and risk.

¹⁴ In order to keep matters simple, I am here focusing on risk as it applies to events, in keeping with the event-focused notion of luck offered earlier. In both cases, however, we can extend these accounts to targets other than events. (Indeed, I think the more natural focal point for risk ascriptions is often *activities* rather than events).

rather than vice versa. Risk tends to concern hazardous events, such as being shot, and hence events that we wish to avoid. Finding out that we are lucky not to have been shot is alarming precisely because it entails that we were at a high risk of being shot.

The second way that the notions of luck and risk diverge is that assessments of risk are always made relative to a specific risk event (or risk events). When evaluating the risk involved in, say, flying a plane, we can do so relative to a range of target risk events, such as the plane crashing, one having a panic attack mid-flight, being sat next to a noisy passenger, and so on. Which risk event we focus on will have a bearing on whether the activity counts as risky. The prospect of the plane crashing may be low risk, for example, but the prospect of being seated next to a noisy passenger may well be high risk.

In contrast, assessments of luck, as we noted above, don't focus on specific events in this way, but rather concern the non-obtaining of the target event, given the same relevant initial conditions, across close possible worlds. This is one important way that the modal account of risk diverges from the modal account of luck, in that what we are interested in is how modally close the risk event is, given that we keep the initial conditions fixed. So, for example, if the risk event associated with taking a flight is the plane crashing, then the question we want to settle is how modally close is the scenario where all the relevant initial conditions are kept fixed (i.e., the subject continues to board the plane, the plane is as well-constructed as it is in the actual world, and so on) but the subject crashes. If this possible world is close, then this activity of taking the flight is high risk, but if it is not, then it is low risk. As before, we also have a continuum of riskiness in play here, depending on how modally close, or far-out, the target risk event is.¹⁵ The crucial difference with the modal account of luck, however, is that the modal account of risk is specifically concerned with the modal closeness of a particular risk event, rather than with the general non-obtaining of the target event in the modal neighbourhood.¹⁶ See also the application of this notion of risk to particular domains, such as aesthetics and law, in [Pritchard](#) (forthcoming, 2016b, 2017b). For a recent account of risk which overlaps in some respect with the one just offered, but which also diverges in some important respects too, see [Broncano-Berrocal](#) (2015).

Putting these two points together, how will an anti-risk epistemology differ from an anti-luck epistemology? Let us begin with the lacunae in anti-luck epistemology that anti-risk epistemology enables us to evade. These lacunae have not been spotted, but that is because they are so well hidden. The first concerns the way in which we have transplanted our modal account of luck into an anti-luck epistemology.

Recall that the luckiness of an event depends on how 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 where it is an easy possibility that one forms a false belief. 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. The question is thus why anti-luck epistemology opts

¹⁵ Though bearing in mind, when it comes to this continuum, that luck is much more of a threshold notion than risk is see footnote 13 for further discussion of this point.

¹⁶ For further discussion of this notion of risk, see [Pritchard](#) (2015b).

for the more restrictive rendering of the anti-luck condition (i.e., where the relevant sense of cognitive failure is understood as false belief specifically, 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. In particular, it would rule out as knowledge those cases involving evidential epistemic luck, even though we saw above that this species of epistemic luck is entirely compatible with 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 translates as either non-belief or false belief), then one would be obliged to treat the agent as lacking knowledge, contrary to intuition.

It is thus very important that anti-luck epistemology opts for the restrictive rendering of the 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 the latter 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).¹⁷ 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.¹⁸

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 from one's knowledge, this will translate as the claim that one is looking for a true belief on a particular basis, where the risk event of forming a false belief on that same basis 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.

Anti-risk epistemology is also better placed than anti-luck epistemology to explain why any belief formed in a necessary truth or even, for that matter, a proposition whose truth is modally stable across close possible worlds, such that there is no close possible world where it is false is not by default a safe true belief, regardless of how that belief was formed. There is, after all, no close possible world where the target belief in question is false. Anti-luck epistemology deals with this problem by appealing to the basis-relativity of safety, something which we saw is motivated by the anti-luck epistemology methodology. The thinking is that we only need to keep the actual basis for belief fixed when we are examining the modal environment, and should not also keep fixed the particular propositional content that is believed in the actual world. Instead, we should be interested in cognitive failure that results from that basis for belief in close possible worlds even when it concerns a different propositional content. In this way, for example, while there is no close possible world where forming a true

¹⁷ This example is an adaption of Radford (1966) 'diffident schoolboy' case, albeit to illustrate a rather different point.

¹⁸ There are lots of other putative counterexamples to safety, of course. I deal with such cases in various places, but see, for example, (Pritchard 2012b, 2015a).

belief in a mathematical proposition via a cognitive process that is lacking in any epistemic pedigree (flipping coins, say) will lead to a belief *in that same proposition* which is false, there will be close possible worlds where that cognitive process will generate false beliefs in related propositions (e.g., where one is led to believing that ‘ $2 + 2 = 5$ ’ rather than ‘ $2 + 2 = 4$ ’)¹⁹ for further discussion of how anti-luck epistemology deals with safe beliefs in necessary/modally stable truths.

The solution is elegant, but it isn’t at all clear that it is strictly speaking licensed by anti-luck epistemology. The focus on the basis is principled, for sure, but why does it follow from the fact that we need to keep the basis fixed that we can vary the doxastic output from this basis across close possible worlds (aside, of course, from the fact that doing so allows us to solve this problem with necessary truths)?

As before, the shift to an anti-risk epistemology resolves the problem, and via the same means. For notice that by putting our focus on a specific risk event of cognitive error we both in a sense narrow and in a sense broaden the ambit of our epistemological concerns. It was a narrowing of our ambit of concern that dealt with the former lacuna in anti-luck epistemology, by making us focus on cognitive error specifically rather than the mere non-obtaining of the target epistemic event of true belief. But given that our spotlight is now on how a certain basis might lead to cognitive error, it also becomes clear why we need to broaden our ambit also to take in cases where this same basis leads to cognitive error in close possible worlds where the doxastic output of the basis is different from that in the actual world. That is, our concern is whether a basis for belief that in the actual world led to a belief that p generates cognitive error in close possible worlds, and not the narrower concern of whether a basis for belief that in the actual world led to a belief that p generates *false beliefs that p* in close possible worlds. We thus get the principled case we require for dealing with safe beliefs in necessary/modally stable truths that was lacking in anti-luck epistemology.

In virtue of its focus on a specific epistemic risk event, anti-risk epistemology is also more easily extendable than anti-luck epistemology to other epistemic goods besides knowledge. The relevant epistemic risk event when it comes to knowledge is cognitive error (on the same basis), but we can easily conceive of all kinds of scenarios where the relevant epistemic goals are such that a distinct epistemic risk event (or events) will be relevant. For example, we can imagine certain kinds of inquiry where the failure to form true beliefs might be a very important epistemic risk event, and so may be just as important indeed, may even be more important than the epistemic risk event of cognitive error.²⁰ More generally, just as our assessments of risk can vary substantially depending on which risk event (or events) we are concerned with, so we can imagine a whole range of epistemic assessments which are structured in terms of epistemic risk events that are different in various ways from the specific epistemic risk event at issue when it comes to knowledge (and which may involve the inherent complexity that arises when there is more than one epistemic risk event in play).²¹

¹⁹ See Pritchard (2007, 2015a).

²⁰ For example, one can imagine certain kinds of scientific inquiry where the pay-off for success is so great that missing out on true beliefs carries a particularly high premium.

²¹ These aren’t the only advantages that anti-risk epistemology has over anti-luck epistemology. For example, the former is better able to account for our intuitions in lottery-style cases than anti-luck epistemology.

3 Anti-risk epistemology and negative epistemic dependence

We noted earlier that when it comes to events that are evaluated in terms of risk and luck, the concern with luck is often in an important sense parasitic on the concern with risk. It is significant that one was lucky not to be shot precisely because it reveals that one was at a high risk of being shot, something that any sensible person would clearly want to avoid. The same goes for epistemic events that are evaluated in terms of risk and luck. We might say of the same true belief that, given how it was formed, it is both only true as a matter of luck and that it was at a high risk of being false. But we are only really concerned with the former because of the latter. Put another way, our concern with eliminating veritic epistemic luck from knowledge is motivated by our desire to eliminate the corresponding veritic epistemic risk i.e., the high risk, given how our belief was formed, that it would have led us into cognitive error.

This point is important, since it helps us to realise just why we want to eliminate veritic epistemic luck from knowledge. One might think that this point would be redundant, given that most epistemologists have granted that knowledge excludes this kind of epistemic luck. In the recent literature, however, there has been a growing wave of commentators who have argued that knowledge can at least in certain cases be compatible with veritic epistemic luck.

In order to understand where this line of thought comes from, we need to distinguish between two kinds of cases where veritic epistemic luck is involved. The first kind of case is of the standard Gettier-style sort, whereby something intervenes between the subject's justified believing and her cognitive success to ensure that the latter obtains in a way that is completely independent of the former. Consider Roderick Chisholm's (1977, p. 105) Gettier-style case involving the farmer who sees a sheep-shaped object in excellent environmental conditions, and who believes truly as a result that there is a sheep in the field before him, but who is in fact looking at a sheep-shaped object (which is obscuring from view the genuine sheep behind). Here we have veritic epistemic luck, and also veritic epistemic risk, in that given how the belief was formed it is both lucky that the belief was true (it could have easily been false) and there was a high risk that the belief was false. Call this kind of veritic epistemic luck *intervening epistemic luck* (where there is also a corresponding kind of veritic epistemic risk that is *intervening epistemic risk*).

Intervening epistemic luck is to be contrasted with a second kind of veritic epistemic luck where the luck in play does not concern anything intervening between justified belief and cognitive success, but rather merely factors in the environment. The famous barn façade example is a case in point.²² For recall that the agent in this case, unlike our agent in Chisholm's sheep case, really is seeing a genuine barn in conditions that would normally be perfectly fine for spotting a barn. The problem, however, is that she is in an environment where almost everything that looks like a barn is in fact a

Footnote 21 continued

Anti-risk epistemology is also able to offer a better diagnostic story regarding why the post-Gettier epistemological project took so many false turns. But these are issues that will need to be explored on another occasion. For further discussion of the merits of anti-risk epistemology, see Pritchard (2016b).

²² As originally proposed in Goldman (1976), but which is credited in turn to Carl Ginet.

barn façade, and hence although her belief is true it is nonetheless subject to veritic epistemic luck (and, for that matter, veritic epistemic risk). Call this kind of veritic epistemic luck/risk *environmental epistemic luck* (where there is also a corresponding kind of veritic epistemic risk that is *environmental epistemic risk*).²³

Although environmental epistemic luck and intervening epistemic luck are both cases of veritic epistemic luck, and hence should both be incompatible with knowledge, this has been called into question by some commentators. In particular, they have argued that environmental epistemic luck might be compatible with knowledge. The rationale for this is that there is a kind of positive epistemic status in play with regard to cases of environmental epistemic luck that isn't applicable to intervening epistemic luck. In particular, cases involving environmental epistemic luck specifically tend to involve not just a justified true belief, but also a *cognitive achievement* i.e., a cognitive success that is properly attributable to one's manifestation of relevant cognitive agency.

We can see this point in action by considering achievements more generally i.e., successes that are properly attributable to one's manifestation of relevant agency. The point is that intervening luck undermines genuine achievements. Suppose one skillfully fires an arrow at a target, and hits the target, but only does so because a dog leaps up mid-flight, seizes the bolt, and inserts it into the target. This is a case where one's relevant agency is manifested (one's archery skills), and also where one is successful, but where there is not the right kind of explanatory relationship between the former and the latter. One is rather successful because of the intervention of the dog rather than because of one's manifestation of relevant agency.

In contrast, imagine a case where nothing in fact intervenes between one's manifestation of one's archery ability and one's successfully hitting the target, but where something could have easily intervened. Suppose, for example, that the dog very nearly caught the bolt mid-flight, but was momentarily distracted and so happened to miss it on this specific occasion. We now have a case involving purely environmental luck. Just as in the previous example, it is a matter of luck that one is successful, in that one could very easily have been unsuccessful (Relatedly, in both cases there is a high risk of failure). But in the second case involving environmental luck, however, the explanatory relationship between one's success and one's manifestation of relevant agency does not seem to be undermined. As we might naturally put it, it is because of one's excellent archery abilities that one hits target, even though the pesky dog might well have easily intervened. Accordingly, it is no less of an achievement.

This point is significant once we transplant it back to the epistemic realm. This is because by parity of reasoning we should say that although intervening epistemic luck precludes the possibility of a cognitive achievement, cases of environmental epistemic luck, in contrast, are entirely compatible with there being a genuine cognitive achievement in play (i.e., where the cognitive success is properly attributable to the subject's manifestation of relevant cognitive agency). But if that's right, then there is something epistemically better about the subject's belief in cases of environmental epistemic luck than in cases of intervening epistemic luck. Accordingly, given this

²³ I originally introduced the distinction between intervening and environmental epistemic luck in Pritchard (2009b), chs. 3–4). For further discussion of this distinction, see Pritchard (2009c, 2012) and Pritchard et al. (2010), chs. 2–4.

epistemic disanalogy between the two cases, what is to prevent one from ascribing knowledge to the agent in the former case? In particular, the temptation is to argue that although the intervening variety of veritic epistemic luck is incompatible with knowledge, the kind of environmental veritic epistemic luck found in barn façade-style case is in fact epistemically benign.²⁴

Elsewhere I have expressed the issue raised by environmental epistemic luck in terms of a particular kind of *epistemic dependence*. By epistemic dependence I mean the ways in which whether one's cognitive success can amount to knowledge can significantly depend upon factors outwith one's cognitive agency.²⁵ I claim that epistemic dependence of this kind comes in both a *negative* and a *positive* form. The latter is when one exhibits a level of cognitive agency that would ordinarily not suffice for knowledge, and yet one knows nonetheless due to a contribution from factors outwith one's cognitive agency.²⁶ The kind of epistemic dependence that presently concerns us, however, is *negative epistemic dependence*. This is when one exhibits a level of cognitive agency that would normally suffice for knowledge as would certainly be the case where one exhibits a cognitive achievement but where one fails to know nonetheless due to factors outwith one's cognitive agency that ensure that one's cognitive success could very easily have been a failure. Instances of environmental epistemic luck are a case in point, in that this seems to undermine one's knowledge even though one's cognitive success is properly attributable to one's manifestation of relevant cognitive agency (and hence qualifies as a cognitive achievement).

The challenge, however, is to explain why such cases of negative epistemic dependence do not count as genuine instances of knowledge. I think that with an anti-risk epistemology in play we are well placed to offer the required explanation. For what is distinctive of cases of negative epistemic dependence is not just that they involve veritic epistemic luck, but more specifically that they involve veritic epistemic risk. That is, given how the belief was formed, then there is a high risk that one is led into cognitive error (as would have occurred, for example, if the agent in the barn façade case had happened to have been looking at the frontage of one of the fake barns in the vicinity). The point is that we care about eliminating veritic epistemic luck from knowledge because we care about eliminating veritic epistemic risk, rather than vice versa. Once that is made explicit, then we can see that what is problematic about cases of negative epistemic dependence is not the presence of veritic epistemic luck, specifically, but rather that they involve high levels of epistemic risk (and hence veritic epistemic risk).

²⁴ See Sosa (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 (2009).

²⁵ For further discussion of this notion of epistemic dependence, see Pritchard (2016a). See also Kallestrup and Pritchard (2013).

²⁶ I claim that positive epistemic dependence is the best way to make sense of why one can gain testimonial knowledge in friendly epistemic environments by, for the most part (though never exclusively), trusting an informant. For more on this point, and on negative epistemic dependence in general, see Kallestrup and Pritchard (2012) and Pritchard (2016a).

Notice too that this point will be similarly effective against cases of negative epistemic dependence that are purely modal in form. In standard barn façade-style cases the negative epistemic dependence concerns an actual feature of the environment, such as that there are fake barns in the vicinity. But in fact one can construct cases of negative epistemic dependence where the knowledge-undermining feature purely concerns not the subject's actual environment but only her purely modal environment.

We can illustrate this point by appeal to what is known as an *epistemic twin earth* case²⁷. We are to imagine two counterpart agents, one on earth and one on twin earth. These agents are microphysical duplicates of each other, with identical causal histories. The 'local' environment that they are presently causally interacting with is also identical in every respect. Finally, 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 her 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 her modal environment the agent on twin earth is forming her 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 her belief in such a way that it is subject to veritic epistemic luck, 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, then we should be inclined to treat the agent's belief on twin earth as an instance of negative epistemic dependence and hence deny knowledge.

Insofar as one is tempted to argue that standard barn façade-style cases might be *bona fide* cases of knowledge, however, then this temptation is bound to be even stronger in epistemic twin earth cases. For if the putative negative epistemic dependence exclusively concerns a feature of the modal environment, then why think that it should be knowledge-undermining at all? But we now have a rationale for resisting this line of argument, one which directly extends from our case for thinking that standard barn façade-style cases cannot be instances of knowledge. This is that the agent on twin earth is forming her belief in a way that is subject to a high degree of epistemic risk, since the target risk event of forming a false belief via this process is in fact modally very close. As before, whereas it might seem at least debatable whether knowledge is compatible with certain kinds of veritic epistemic luck, it is hard to countenance the idea that knowledge might be compatible with high levels

²⁷ See [Kallestrup and Pritchard \(2014\)](#).

of epistemic risk (i.e., that knowledge is compatible with veritic epistemic risk). But if that's right, then we have a genuine case of negative epistemic dependence, even though the knowledge-undermining factors at issue are purely modal.²⁸

4 Concluding remarks

I want to conclude by emphasizing one important implication of this last point about how anti-risk epistemology underwrites the intuition that knowledge is lacking in epistemic twin earth cases where the epistemic luck in play is purely modal. This is significant in that it reinforces the idea that the project of trying to exclusively understand knowledge in terms of the manifestation of cognitive agency is doomed to failure. Such a project which I have elsewhere termed *robust virtue epistemology* has been defended in various forms by such figures as Sosa (1988, 1991, 2007, 2009), Zagzebski (1996, 1999), and Greco (e.g., 2009). The kind of negative epistemic dependence at issue in epistemic twin earth cases is particularly problematic for robust virtue epistemology. This is because in such cases we are keeping fixed across both earth and twin earth anything that could be remotely relevant to the manifestation of cognitive agency. And yet the agent on twin earth is nonetheless lacking knowledge due to some feature of the purely modal environment. It follows that knowledge simply cannot just be a function of the manifestation of cognitive agency as robust virtue epistemology wants to claim. Instead, we are going to need to have a distinct way of accommodating the fact that knowledge is incompatible with high levels of epistemic risk. This point strongly suggests though it doesn't yet entail that what we should be looking for when it comes to the theory of knowledge is not a robust virtue epistemology but rather the kind of *anti-risk virtue epistemology* that is able to blend both a virtue-theoretic and an anti-risk requirement on knowledge.²⁹ See also (Pritchard et al. 2010, chs. 1–4) and Pritchard (2017a). The failure to explain away cases of negative epistemic dependence thus undermines one prominent account of knowledge and strongly suggests an alternative proposal. But this is a further issue to be explored on another occasion.

Acknowledgements Thanks to Jesus Navarro for helpful discussions on this topic. Thanks also to two anonymous referees for *Synthese* for their comments on the paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

²⁸ For further discussion of cases involving purely modal veritic epistemic luck, see Pritchard (2015a, §3).

²⁹ I say that this proposal is merely suggested rather than entailed, since for all I have said there might be scope for several other kinds of theories of knowledge, including one cast exclusively in terms of an anti-risk epistemic condition. Still, I think that the prospects for developing an anti-risk virtue epistemology are very good, particularly since it would inherit many of the advantages of its sister position, *anti-luck virtue epistemology*, while also not being subject to some of the problems facing the anti-luck element of this view that we have noted here. For more on the rationale for anti-luck virtue epistemology, see especially Pritchard (2012).

References

- Becker, K. (2007). *Epistemology modalized*. London: Routledge.
- Black, T. (2002). A moorean response to brain-in-a-vat scepticism. *Australasian Journal of Philosophy*, 80, 148–163.
- Black, T. (2008). Defending a sensitive neo-Moorean invariantism. In V. F. Hendricks & D. H. Pritchard (Eds.), *New waves in epistemology* (pp. 8–27). Basingstoke: Palgrave Macmillan.
- Black, T. (2011). Modal and anti-luck epistemology. In S. Bernecker & D. H. Pritchard (Eds.), *Routledge companion to epistemology* (pp. 187–198). London: Routledge.
- Black, T., & Murphy, P. (2007). In defense of sensitivity. *Synthese*, 154, 53–71.
- 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.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–23.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–91.
- Greco, J. (2009). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
- Hansson, S. O. (2004). Philosophical perspectives on risk. *Techné: Research in Philosophy and Technology*, 8, 10–35.
- Hansson, S. O. (2014). Risk. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/spr2014/entries/risk/>.
- Hetherington, S. (2013). There can be lucky knowledge. In M. Steup & J. Turri (Eds.), *Contemporary debates in epistemology*, S7 (2nd ed.). Oxford: Blackwell.
- Kahneman, D., & Varey, C. A. (1990). Propensities and counterfactuals: The loser that almost won. *Journal of Personality and Social Psychology*, 59, 1101–10.
- Kallestrup, J., & Pritchard, D. H. (2012). Robust virtue epistemology and epistemic anti-individualism. *Pacific Philosophical Quarterly*, 93, 84–103.
- Kallestrup, J., & Pritchard, D. H. (2013). Robust virtue epistemology and epistemic dependence. In T. Henning & D. Schweikard (Eds.), *Knowledge, virtue and action: Putting epistemic virtues to work*, ch. 11. London: Routledge.
- Kallestrup, J., & Pritchard, D. H. (2014). Virtue epistemology and epistemic twin earth. *European Journal of Philosophy*, 22, 335–57.
- Levy, N. (2011). *Hard luck: How luck undermines free will and moral responsibility*. Oxford: Oxford University Press.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1987). *On the plurality of worlds*. Oxford: Blackwell.
- Luper, S. (1984). The epistemic predicament. *Australasian Journal of Philosophy*, 62, 26–50.
- Luper, S. (2003). Indiscernibility skepticism. In S. Luper (Ed.), *The skeptics: Contemporary essays* (pp. 183–202). Aldershot: Ashgate.
- Mumpower, J. (1986). An analysis of the de minimis strategy for risk management. *Risk Analysis*, 6, 437–46.
- Nozick, R. (1981). *Philosophical explanations*. Oxford: Oxford University Press.
- Petersen, M. (2002). What is a de minimis risk? *Risk Management*, 4, 47–55.
- Pritchard, D. H. (2002). Resurrecting the Moorean response to the sceptic. *International Journal of Philosophical Studies*, 10, 283–307.
- Pritchard, D. H. (2004). Epistemic luck. *Journal of Philosophical Research*, 29, 193–222.
- Pritchard, D. H. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pritchard, D. H. (2006). Moral and epistemic luck. *Metaphilosophy*, 37, 1–25.
- Pritchard, D. H. (2007). Anti-luck epistemology. *Synthese*, 158, 277–97.
- Pritchard, D. H. (2008). Sensitivity, safety, and anti-luck epistemology. In J. Greco (Ed.), *Oxford handbook of scepticism* (pp. 437–455). Oxford: Oxford University Press.
- Pritchard, D. H. (2009a). Apt performance and epistemic value. *Philosophical Studies*, 143, 407–16.
- Pritchard, D. H. (2009b). *Knowledge*. London: Palgrave Macmillan.
- Pritchard, D. H. (2009c). Knowledge, understanding and epistemic value. In A. O’Hear (Ed.), *Epistemology (Royal Institute of Philosophy Lectures)* (pp. 19–43). Cambridge: Cambridge University Press.

- Pritchard, D. H. (2012a). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–79.
- Pritchard, D. H. (2012b). In defence of modest anti-luck epistemology. In T. Black & K. Becker (Eds.), *The sensitivity principle in epistemology* (pp. 173–192). Cambridge: Cambridge University Press.
- Pritchard, D. H. (2013). There cannot be lucky knowledge. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 152–164). Oxford: Blackwell.
- Pritchard, D. H. (2014). The modal account of luck. *Metaphilosophy*, 45, 594–619.
- Pritchard, D. H. (2015a). Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172, 93–111.
- Pritchard, D. H. (2015b). Risk. *Metaphilosophy*, 46, 436–61.
- Pritchard, D. H. (2016a). Epistemic dependence. *Philosophical Perspectives*, 30, 1–20.
- Pritchard, D. H. (2016b). Epistemic risk. *Journal of Philosophy*, 113, 550–71.
- Pritchard, D. H. (2017a). Knowledge, luck and virtue: Resolving the Gettier problem. In C. Almeida, P. Klein, & R. Borges (Eds.), *The Gettier Problem* (pp. 57–73). Oxford University Press.
- Pritchard, D. H. (2017b). Legal risk, legal evidence, and the arithmetic of criminal justice. *Jurispudence*. doi:[10.1080/20403313.2017.1352323](https://doi.org/10.1080/20403313.2017.1352323).
- Pritchard, D. H. (forthcoming). Aesthetic Risk. In *Think*.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The nature and value of knowledge: Three investigations*. Oxford: Oxford University Press.
- Pritchard, D. H., & Smith, M. (2004). The psychology and philosophy of luck. *New Ideas in Psychology*, 22, 1–28.
- Pritchard, D. H., & Whittington, L. J. (Eds.). (2015). *The philosophy of luck*. Oxford: Blackwell.
- Radford, C. (1966). Knowledge by examples. *Analysis*, 27, 1–11.
- Roush, S. (2005). *Tracking truth: Knowledge, Evidence and science*. Oxford: Oxford University Press.
- Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57, 907–19.
- Sosa, E. (1988). Beyond skepticism, to the best of our knowledge. *Mind*, 97, 153–89.
- Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge: Cambridge University Press.
- Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–54.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge*. Oxford: Oxford University Press.
- Sosa, E. (2009). *Reflective knowledge: Apt belief and reflective knowledge*. Oxford: Oxford 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.
- Teigen, K. H. (1996). Luck: The art of a near miss. *Scandinavian Journal of Psychology*, 37, 156–71.
- Teigen, K. H. (1997). Luck envy, gratitude: It could have been different. *Scandinavian Journal of Psychology*, 38, 318–23.
- Teigen, K. H. (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, 235–48.
- Teigen, K. H. (1998b). When the unreal if more likely than the real: Post hoc probability judgements and counterfactual closeness. *Thinking and Reasoning*, 4, 147–77.
- Teigen, K. H. (2003). When a small difference makes a large difference: Counterfactual thinking and luck. In D. R. Mandel, D. Hilton, & P. Catellani (Eds.), *The psychology of counterfactual thinking*. London: Routledge.
- 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, 639–52.
- Tetlock, P. E., & Lebow, R. N. (2001). Poking counterfactual holes in covering laws: Cognitive styles and historical reasoning. *American Political Science Review*, 95, 829–43.
- 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: Cambridge University Press.
- Zagzebski, L. (1999). What is knowledge? In J. Greco & E. Sosa (Eds.), *Blackwell guide to epistemology* (pp. 92–116). Oxford: Blackwell.