

Luck and normative achievements: Let not safety be our guide

Bruno Guindon

Penultimate version - please cite published version when possible

Abstract

It is a well-worn platitude that knowledge excludes luck. According to anti-luck virtue epistemology, making good on the anti-luck platitude requires an explicit anti-luck condition along the lines of *safety*: S knows that *p* only if S's true belief that *p* could not have easily been mistaken. This paper offers an independent, virtue epistemological argument against the claim that safety is a necessary condition on knowledge, one that adequately captures the anti-luck platitude. The argument proceeds by way of analogy. I focus on two paradigmatic kinds of normative achievements that also exclude luck: (i) – having a doxastically justified belief, and (ii) – performing a morally worthy action. I then show that while (i) and (ii) exclude luck, they are nevertheless susceptible to what I call 'modal luck'. I then move on to show that knowledge, or at least some instance of knowledge, is a normative achievement, which I claim provides strong reasons to expect that knowledge is also susceptible to modal luck. Since safety entails that knowledge is incompatible with modal luck, the argument provides strong reasons to reject safety as a necessary condition on knowledge.

1. Introduction

Virtue epistemologists of all stripes agree that S knows that *p* only if

ACHIEVEMENT: S's cognitive success (i.e., S's true belief that *p*) is due to S's cognitive abilities (i.e., S's epistemic virtues).¹

Within the virtue epistemology literature, ACHIEVEMENT is taken to be a necessary condition for knowledge. But there is a live debate whether ACHIEVEMENT is also sufficient. Proponents of robust virtue epistemology (RVE) claim that it is, whereas proponents of modest virtue

¹ While every virtue epistemologist should agree with this way of putting the view, there is significant debate about how to understand the 'due to' relation. Compare, e.g., (Carter, 2013; Greco, 2009; Jarvis, 2013; Sosa, 2007; Turri, 2011; Zagzebski, 1996). I hope to remain as neutral as possible about the 'due to' relation here.

epistemology (MVE) deny it.² Why do proponents of MVE deny that ACHIEVEMENT is sufficient for knowledge? They appeal to the platitude that knowledge excludes various kinds of luck and argue that ACHIEVEMENT is compatible with knowledge-undermining luck.³

Consider Alvin Goldman's famous barn façade case:

FAKE BARNS: Henry is driving through the countryside, and unbeknownst to him, he has just entered Fake Barn County, a district filled with papier-mâché barn facsimilia. As it happens, Henry is looking at the only real barn in the county and forms the belief that he is looking at a barn.⁴

According to proponents of MVE, while it may be granted that Henry's cognitive success is due to his cognitive abilities, there is an overwhelming intuition in FAKE BARNS that Henry does not know that he is looking at a barn.⁵ The reason that is offered for why Henry does not know in FAKE BARNS is that his success is also in an important sense a matter of luck: given the abundance of barn façades in the neighborhood, Henry's cognitive abilities could have easily ended in cognitive failure (i.e., false belief). Proponents of MVE take this to show that capturing the anti-luck platitude requires an independent anti-luck condition on knowledge, something along the lines of the following modal principle:

SAFETY: S knows that *p* only if S's belief that *p* could not have easily been false.⁶

² For proponents of RVE, see (Carter, 2013; Greco, 2003, 2012; Jarvis, 2013; Sosa, 2007; Zagzebski, 1996). See also (Littlejohn, 2014), who is sympathetic to the view. For proponents of MVE, see (Kelp, 2013; Pritchard, 2007, 2012a, 2012b, 2015, 2016).

³ Although, see (Jarvis, 2013) and (Carter, 2013) who deny that ACHIEVEMENT is compatible with knowledge-undermining luck.

⁴ See (Goldman, 1976). See also (Goldman, 1992, 2009), who subsequently credits the example to Carl Ginet.

⁵ That Henry's cognitive success is due to his cognitive abilities is not entirely uncontroversial. Again, see (Carter, 2013) and (Jarvis, 2013) who deny it. According to Jarvis, while Henry's cognitive abilities are indeed exercised, they are nevertheless not *manifested*, and thus not responsible for his cognitive success. For Carter, Henry's success is not due to his cognitive abilities because his success is not *primarily creditable* to his abilities. Pritchard (2016), on the other hand, thinks that there is a case to be made that Henry's cognitive success is due (or 'attributable', in Pritchard's terms) to his cognitive abilities (p. 569). Since my main opponent is the proponent of MVE, and Pritchard is one of its main defenders, I will follow suit and grant that Henry's cognitive success is due to his cognitive abilities.

⁶ I write 'something along the lines of SAFETY' since no one actually endorses SAFETY in its current form. The principle needs to be refined to avoid obvious objections. Following Pritchard (who is its most ardent defender), we

Yet, there is an air of circularity to the argument. Intuitions about the lack of knowledge in cases like FAKE BARNS and the idea that SAFETY (or something like it) adequately captures the anti-luck platitude are mutually self-supporting. Why does SAFETY adequately capture the anti-luck platitude? We are told that it is because it rightly predicts a lack of knowledge in cases like FAKE BARNS. But if we ask why there is a lack of knowledge in such cases, we are told that it is because the beliefs in question could have easily been false. So, the answer presupposes SAFETY.

This paper offers an independent, virtue epistemological argument against SAFETY (or something like it) as a necessary condition on knowledge, one that adequately captures the anti-luck platitude. Some have argued against safety principles by showing that there are cases of unsafe knowledge—i.e., cases of knowledge where an agent’s belief could have easily been false.⁷ I take these to be genuine counterexamples to SAFETY. But the argument in this paper does not essentially rely on examples of unsafe knowledge. Rather, it argues by way of analogy. After laying out some preliminaries in section 2, I consider in section 3 two distinct types of normative achievements that also exclude luck: (i) – having a doxastically justified belief, and (ii) – performing a morally worthy action. I argue here that (i) and (ii) are nevertheless susceptible to what I call ‘modal luck’, and thus unconstrained by SAFETY-like modal principles.

In section 4, I argue that, at least for some instances of knowledge, knowledge is also a kind of normative achievement. Of course, it does not strictly follow from this that knowledge is *thereby* susceptible to modal luck. It is possible that only some but not all normative

might call the view that takes SAFETY and ACHIEVEMENT to be both necessary and sufficient for knowledge ‘anti-luck virtue epistemology’. Cf. (Kelp, 2013), who defends a similar ‘safe-apt’ view of knowledge.

⁷ See, e.g., (Comesaña, 2005; Kelp, 2009) who argue that there can be unsafe knowledge, although Kelp (2013) has since changed his view. See also (McBride, 2011) who argues that there can be unsafe legal knowledge.

achievements are so susceptible. Still, short of a principled reason why we should expect knowledge to be different in this regard, anyone who thinks that ACHIEVEMENT is a necessary condition on knowledge should take the arguments from sections 3 and 4 to offer a compelling case that knowledge, or at least some instance of knowledge, is also susceptible to modal luck. Since SAFETY excludes modal luck, the argument provides independent grounds for rejecting SAFETY as an anti-luck condition on knowledge, and thus undermines the need to defend a more modest form of virtue epistemology. Section 5 concludes by considering some possible implications my argument has on the knowledge-undermining status of environmental epistemic luck.

2. The setup

It is a well-worn platitude in epistemology that knowledge excludes luck: If S knows that p , then S truly believes that p , and it is not a matter of luck that S truly believes that p .^{8,9} Yet, as Pritchard (2005, 2007) makes clear, not every kind of luck is knowledge-undermining. As a matter of luck, I just so happen to overhear the commis inform the chef de partie that the bisque has gone bad. Subsequently, I form the belief that the bisque has gone bad. I am lucky (indeed, quite fortunate) to have a true belief about the state of the bisque: I now know not to order it. But

⁸ This is what (Gettier, 1963), and the subsequent extensive Gettier-ology literature has been taken to have shown. For an early discussion of knowledge-undermining luck, see (Russell, 2013), although Russell thought that epistemic luck, when knowledge-undermining, was so in virtue of undercutting justification. For an even earlier treatment, see (Krasser, 1995), who finds Gettier-style cases in a text dated from 770 CE by Indian philosopher Dharmottara.

⁹ I am here assuming that knowledge entails belief. I take this to be fairly uncontroversial, although some have sought to rejected it. See, e.g., (Radford, 1966). I do not share many of the intuitions that Radford marshals in favor of his argument, but my sense is that if one shares these intuitions, they can be explained away by drawing a distinction between knowledge and factual awareness. For an interesting and persuasive defense of the distinction between knowledge and factual awareness, see (Silva, 2023).

this kind of luck—i.e., luck in the evidence one possesses—does not undermine the fact that I, unlike most dinner guests, know that the bisque has gone bad.¹⁰

One kind of epistemic luck that is widely understood to be knowledge-undermining is the kind of luck we find in typical Gettier-style cases, what Pritchard calls ‘intervening epistemic luck’.¹¹ Here is one example due to Roderick Chisholm (1977):

SHEEP: You are standing in a field looking at what appears to you to be a sheep. You subsequently form the belief that there is a sheep in the field. Unfortunately, what you are looking at is a hairy dog that merely looks like a sheep. There is a sheep in the field, however, hiding out of sight behind the hairy dog.

In SHEEP, the overwhelming intuition is that while you might have a justified true belief that there is a sheep in the field, you do not know that there is a sheep in the field. A natural explanation of why you do not know this is that it is a matter of luck that your belief is true: Had there been no sheep in the field, you would still have believed that there is a sheep in the field, and that belief would have been false. It is a lucky accident that a real sheep *intervenes*, as it were, rendering your belief true.¹²

Why is intervening epistemic luck, and not luck in the evidence one possesses, incompatible with knowledge? RVE offers us a good explanation. According to RVE, knowledge is a particular kind of cognitive achievement. It is cognitive success (true belief) due to cognitive abilities (epistemic virtues). In the bisque case, your true belief is due to your

¹⁰ Pritchard calls this ‘evidential epistemic luck’. See (Pritchard, 2005, 2007) for other kinds of benign cases of epistemic luck.

¹¹ See (Pritchard, 2009, 2016) for the distinction between intervening and environmental epistemic luck. More on environmental epistemic luck below.

¹² The argument in this paper does not depend on any substantive view about the nature of luck. Still, I find Pritchard’s modal account to be particularly insightful. Roughly, an event is lucky when it is an event that actually obtains but does not obtain in a wide range of close possible worlds in which the event’s relevant conditions are held fixed. According to this account, there being a real sheep hiding behind the hairy dog is a lucky event since there is a wide range of close possible worlds where the relevant conditions are held fixed but where there is no sheep in the field. For a recent defense of the modal account of luck, see (Pritchard, 2014). See also (Pritchard, 2007, 2012b) for earlier defenses.

cognitive abilities. This is not so in SHEEP. In SHEEP, your cognitive success is due in large part to the lucky fortune that there happens to be a sheep hiding out of sight. Intervening epistemic luck, unlike luck in the evidence one possesses, is incompatible with knowledge because in cases of intervening epistemic luck, your cognitive abilities play no significant role in your cognitive success.

While RVE can explain why intervening epistemic luck is incompatible with knowledge, some virtue epistemologists are skeptical that ACHIEVEMENT can fully capture the anti-luck platitude, and thus serve as the anti-luck condition on knowledge.¹³ In addition to intervening epistemic luck, it is widely assumed that environmental epistemic luck is also knowledge-undermining. This is the kind of epistemic luck we find in FAKE BARNS. It is often claimed that Henry does not know that he is looking at a barn.¹⁴ Again, the typical explanation is that, like in cases of intervening epistemic luck, it is a lucky accident that Henry's belief is true: had Henry been looking at one of the many barn façades around, he would have formed the same belief, yet his belief would have been false. The problem for proponents of RVE, it is argued, is that there is ACHIEVEMENT in FAKE BARNS: Henry has a true belief, and his cognitive success is due to his cognitive abilities.¹⁵ RVE predicts that Henry knows that he is looking at a barn. So, while proponents of RVE can rely on ACHIEVEMENT to explain why intervening epistemic luck is knowledge-undermining, they cannot use it to explain why environmental epistemic luck is equally knowledge-undermining.

¹³ See, e.g., (Kallestrup & Pritchard, 2016; Kelp, 2013; Pritchard, 2007, 2009, 2012b, 2012a, 2015, 2016).

¹⁴ That Henry lacks knowledge in FAKE BARNS is a pervasive assumption in the literature, although it is seldom argued for. For a few examples, see (Goldman, 1976; Pritchard, 2005, 2007; Silva, 2023). Lycan (2006) and Sosa (2007) are two of the few who claim that Henry knows in FAKE BARNS. See also (Lycan, 1977), who argues against the idea that unpossessed evidence, like the fact that Henry is driving through Fake Barn County, can act as undercutting defeaters.

¹⁵ Again, this is not uncontroversial. See fns. 3 and 5 above.

The failure of ACHIEVEMENT in explaining why environmental epistemic luck is knowledge-undermining has prompted some virtue epistemologists to defend more modest versions of virtue epistemology.¹⁶ Proponents of MVE agree that ACHIEVEMENT is necessary for knowledge, but they deny that it is sufficient. According to Pritchard, what is needed to handle cases of environmental luck is an additional modal condition on knowledge, something along the lines of

SAFETY: S knows that *p* only if S's belief that *p* could not have easily been false.

As mentioned earlier, SAFETY is, in its current form, elliptical. It is important that the principle be relativized to the actual basis of the agent's belief that *p*, since it otherwise would be susceptible to obvious counterexamples. Suppose you lack any belief about Anna's astrological sign, but you now take a guess. You now believe that Anna is a Libra. But this is quite unlike you. You typically do not form beliefs about people's astrological signs without asking. Typically, you ask, and when you do, people answer truthfully. But by chance, you guessed correctly. So, here is a case in which your true belief could not have easily been mistaken: had you not formed the belief by merely guessing, you would have asked Anna (as you typically do), and she would have told you truthfully that she is a Libra, causing you to believe truly that she is a Libra. Still, we would not want to count your belief as a candidate for knowledge. Mere guesswork is no path to knowledge.

Once we hold fixed the basis of your belief and focus on instances where you form your beliefs about people's astrological signs via mere guesswork, it becomes obvious that your true belief could have easily been mistaken. Following Pritchard (2009), we can flesh out the principle in terms of possible worlds:

¹⁶ See (Kelp, 2013; Pritchard, 2007, 2012a, 2012b, 2015, 2016).

SAFETY-FLESHED-OUT: S knows that p only if in most nearby possible worlds in which S continues to form her belief that p in the same way as in the actual world, and in all very close nearby possible worlds in which S continues to form her belief that p in the same way as the actual world, p is true.

For those who think that environmental epistemic luck is knowledge-undermining, SAFETY-FLESHED-OUT seems to offer a good explanation why: there are very close nearby possible worlds in which Henry continues to believe that he is looking at a barn, where his belief is based on his perceptual observation, but where his belief is false.¹⁷ Those are the very nearby possible worlds in which Henry is not looking at a barn but at one of the many barn façade facsimilia in the county.

But is the kind of luck involved in FAKE BARNS knowledge-undermining? The assumption that environmental epistemic luck is knowledge-undermining is widespread, and it plays an important role in various arguments against RVE. For example, it is a crucial premise in Kallestrup & Pritchard's (2014) Epistemic Twin Earth argument against RVE. It also serves as a vital assumption in Kallestrup & Pritchard's (2013) argument for what they call the “epistemic dependence thesis”, which, they argue, is incompatible with RVE. But while it is often assumed that environmental epistemic luck is knowledge-undermining, the claim is seldom argued for. Granted, many who reflect on cases of environmental epistemic luck have strong intuitions that the protagonists in such cases lack knowledge. And yet, there are reasons for being skeptical.

One reason for thinking that environmental epistemic luck is knowledge-undermining might be grounded in the conviction that what we mean when we say that knowledge excludes luck is that, when one knows, one's belief could not have easily been false. This, of course, is just to say that SAFETY adequately captures the anti-luck platitude on knowledge. Yet, if we ask, “Why does SAFETY adequately capture the anti-luck platitude on knowledge? Why does

¹⁷ In the remainder of this paper, I will use ‘SAFETY’ as a short-hand for ‘SAFETY-FLESHED-OUT’.

ACHIEVEMENT not adequately capture it?”), we are told that it is because SAFETY can explain why there is a lack of knowledge in cases like FAKE BARNS, while ACHIEVEMENT alone cannot. So, the answer begs the question: it assumes that environmental epistemic luck is knowledge-undermining.

Secondly, while many share the intuition that there is a lack of knowledge in cases like FAKE BARNS, there certainly is no consensus.¹⁸ As Gendler & Hawthorne (2005) have shown, intuitions about the lack of knowledge in cases of environmental epistemic luck are quite unstable and depend a great deal on the details of the cases. Additionally, Colaço et al. (2014) argue that knowledge attributions in cases like FAKE BARNS negatively correlate with the ages of the would-be attributors: the older the would-be attributor, the less likely they are to attribute knowledge in cases like FAKE BARNS. The converse is also true: the younger the would-be attributor, the more likely they are to attribute knowledge to protagonists in cases of environmental epistemic luck. So, some people do attribute knowledge in cases of environmental epistemic luck. The upshot, of course, is not that environmental epistemic luck is compatible with knowledge, after all. Rather, it is that we should not let intuitions settle the question either way.

How should we settle the question? From the perspective of virtue epistemology, the question ‘Is environmental epistemic luck knowledge-undermining?’ is really a question about whether a cognitive achievement like knowledge is possible in environments where a success due to abilities could have easily been a cognitive failure. But here, virtue epistemologists are not without comparative resources. There are other kinds of cognitive achievements, and seeing whether those are constrained by SAFETY-like principles should help settle the issue.

¹⁸ Again, see (Lycan, 1977, 2006; Sosa, 2007).

I begin the following section by focusing on *one instance* of having a doxastically justified belief—i.e., where one’s justified belief is based on the normative reasons that justify it. As we will see, such an instance is a *normative* cognitive achievement, since it is a species of the more general normative achievement of correctly responding to one’s possessed normative reasons. And while not entirely cognitive, I also consider the normative achievement of performing a morally worthy action. I argue that, like knowledge, these normative achievements exclude a certain kind of luck. Nevertheless, I argue that they are susceptible to what I call ‘modal luck’.¹⁹ As I understand it, modal luck is luck in the possible world in which one finds oneself. Specifically, it is luck in being in a possible world in which one’s normative achievement is not undermined. I show that in virtue of being susceptible to modal luck, such normative achievements are not constrained by SAFETY-like principles. So, it is consistent with these normative achievements that they could have easily been normative failures. This will constitute the first part of my argument against the claim that SAFETY is a necessary condition on knowledge, one that adequately captures the anti-luck platitude. The second part will be to show in section 4 that knowledge, or at least some instance of knowledge, is also a normative achievement.

3. Normative achievements and modal luck

3.1. Doxastic justification

¹⁹ For now, I remain agnostic about whether modal luck is reducible to the more familiar forms of epistemic luck we find in the literature. I briefly come back to this issue in section 5.

Epistemologists like to distinguish between *propositional* and *doxastic* justification.²⁰

Propositional justification concerns whether there is sufficient justification to believe a certain proposition. Doxastic justification, on the other hand, concerns whether a certain belief is justifiably held. One important difference between the two is that there can be propositional justification without belief. Given that I know that the bisque has gone bad, I have sufficient reason to believe that eating it will make me ill. So, the proposition <eating the bisque will make me ill> is propositionally justified for me, and this is so even if I do not (perhaps only yet) believe that eating it will make me ill.²¹ Doxastic justification, on the other hand, requires belief *and* propositional justification. Once I form the belief that is propositionally justified for me, we can then ask whether my belief is also doxastically justified.

Is it doxastically justified? Most philosophers agree that having a doxastically justified belief is more than having a belief that is propositionally justified. Suppose I believe that eating the bisque will make me ill, but I come to have this belief by flipping a coin. Surely, I believe what I am propositionally justified in believing, yet my belief is not justifiably held. As most philosophers agree, for my belief to be doxastically justified, there needs to be a connection between my belief that eating the bisque will make me ill and that which makes it

²⁰ Some like to flesh out the distinction in terms of ‘*ex ante*’ and ‘*ex post* justification/rationality’. See, e.g., (Goldman, 2012; Lord, 2018), although see (Jarvis & Ichikawa, 2013) who deny that the propositional/doxastic distinction maps onto the *ex ante/ex post* distinction. The distinction between *ex ante* and *ex post* warrant was first introduced by Firth (1978). The literature on this distinction is vast and includes various debates that I simply cannot address here. For a useful survey of some of the issues, see (Silva & Oliveira, 2023).

²¹ While I think that S having sufficient reason to believe that *p* is sufficient for *p* to be propositionally justified for S, I remain silent on whether it is also necessary. This leaves open the possibility of understanding the notion of propositional justification outside a ‘reasons-first’ framework. Related, one important debate in the literature is whether we ought to understand doxastic justification in terms of ‘propositional justification’ or vice versa. Many are inclined to think the former. See, e.g., (Silva & Oliveira, 2023). Others think that this is backwards. See, e.g., (Goldman, 2012; Turri, 2010). Neta (2022) argues that neither notion is reducible to the other. I remain silent on which notion—propositional or doxastic justification—if any, is more fundamental. All I need for the sake of this paper is that S’s belief that *p* is doxastically justified only if *p* is propositionally justified for S, and that if S has sufficient reason to believe that *p*, then *p* is propositionally justified for S.

propositionally justified.²² But not any old connection will do. Consider a case from John Turri (2010, 315-6):

Consider two [...] jurors, Miss Proper and Miss Improper, sitting in judgment of Mr. Mansour. Each paid close attention throughout the trial. As a result, each knows the following things:

- (P1) Mansour had a motive to kill the victim.
- (P2) Mansour had previously threatened to kill the victim.
- (P3) Multiple eyewitnesses place Mansour at the crime scene.
- (P4) Mansour's fingerprints were all over the murder weapon.

<Mansour is guilty> is propositionally justified for both jurors because each knows (P1 – P4). As it happens, each comes to believe <Mansour is guilty> as the result of an episode of explicit, conscious reasoning that features (P1 – P4) essentially. Miss Proper reasons like so:

(Proper Reasoning) (P1 – P4) make it overwhelmingly likely that Mansour is guilty. (P1 – P4) are true. Therefore, Mansour is guilty.

Miss Improper, by contrast, reasons like this:

(Improper Reasoning) The tea leaves say that (P1 – P4) make it overwhelmingly likely that Mansour is guilty. (P1 – P4) are true. Therefore, Mansour is guilty.

For both Miss Proper and Miss Improper, there is a connection between their beliefs and what makes their beliefs propositionally justified: (P1 – P4) explicitly appear in their reasonings. But in the case of Miss Improper, the connection between (P1 – P4) and her belief is *deviant*, mediated by the tea leaves. So, while Miss Proper's belief that Mansour is guilty is justifiably held, Miss Improper's belief is not.

Here is another deviant case from Errol Lord (2018, 132):

El Clasico

Sam is wondering whether Real Madrid lost yesterday. She believes that Barcelona won yesterday. This is, as it happens, a strong normative reason to

²² See, e.g., (Kvanvig, 2003; Pollock & Cruz, 1999; Swain, 1979).

believe that Real Madrid lost, as yesterday was the El Clasico game between Barcelona and Real Madrid. Sam infers that Real Madrid lost from her belief that Barcelona won. The inference rule she follows is Explosion: Infer <REAL MADRID LOST> from any proposition *p*.²³

Like Miss Improper, there is a connection between Sam's belief and what propositionally justified it: her belief is, *ex hypothesi*, inferred from what propositionally justifies it. Yet, since her belief is generated by some absurd rule of reference, its connection with what propositionally justifies it is deviant.

Cases like that of Miss Improper and El Clasico—cases of what we might call 'epistemic deviancy'—reveal that doxastic justification requires there to be an *appropriate, non-deviant*, connection between an agent's belief and what propositionally justifies it.²⁴ For the kind of doxastic justification that I have in mind—i.e., where S is propositionally justified in believing that *p because* S possesses sufficient reasons to believe that *p*—it will suffice to think about the appropriate connection in terms of 'proper basing', such that:

DOXASTIC JUSTIFICATION (DJ): S's belief that *p* is doxastically justified if (i) S has sufficient reasons R to believe that *p*, (ii) S believes that *p*, and (iii) S properly bases their belief that *p* on R.

²³ If it helps, we can include in Sam's epistemic ken the true proposition that Real Madrid played Barcelona yesterday.

²⁴ This is the main point defended in (Silva, 2015b). Lord (2018) and Lord & Sylvan (2019) take cases of epistemic deviancy to motivate their 'prime view' of reacting for a normative reason, where reacting for a normative reason is not simply a matter of reacting for a motivating (operative) reason that *just so happens* to be a normative reason. Their reasoning, which I take to be quite persuasive, is that any 'composite' account of reacting for a normative reason—accounts holding that reacting for a normative reason is *nothing more* than reacting for a motivating (operative) reason that happens to be a normative reason—is susceptible to cases of epistemic deviancy. I take their view to be compatible with DJ, since they can—and should—precisify the proper basing relation that appears in DJ in terms of 'believing for sufficient normative reasons'. As I will make clear below, a big part of what motivates Lord's and Sylvan's prime view is the same motivation others have in landing on the right account of the proper basing relation for doxastic justification: to avoid cases where it is a lucky accident that one's basis for belief is a sufficient reason to believe it.

A few things to note here. First, DJ is not an analysis or real definition of doxastic justification. DJ only offers a sufficient condition for doxastic justification.²⁵ Still, it is the condition on which I focus in this paper. So, when I speak of ‘doxastic justification’, I mean to speak of only those instances where one’s belief that *p* is doxastically justified *because* one properly bases one’s belief that *p* on the sufficient reasons R that one possesses to believe that *p*.

Second, while DJ explicates doxastic justification in terms ‘proper basing’, it is nevertheless quite neutral. That is because DJ does not commit to any substantive account of proper basing. There is a live debate as to what proper basing amounts to.²⁶ My argument is meant to be ecumenical with respect to this debate. All I need for my argument is what I call ‘Basing Achievement’:

BASING ACHIEVEMENT (BA): If S properly bases their belief that *p* on the sufficient reasons R that S possesses to believe that *p*, then (i) – S’s basing is an exercise of S’s cognitive abilities, (ii) – S is *creditworthy* for exercising said abilities, and (iii) – S being creditworthy for exercising said abilities in basing their belief that *p* on R is what makes S’s basing a *normative achievement*.²⁷

I take it to be uncontroversial that properly basing one’s belief that *p* on the sufficient reasons R that one possesses to believe that *p* is an exercise of one’s cognitive abilities. In any event, I leave this claim as an undefended assumption. And while I do not have the space to provide a full defense of the claim that S’s proper basing is creditworthy and thus a normative achievement, it gets considerable support from reflection on cases of epistemic deviancy like that of Miss Improper and El Clasico.

²⁵ So, it is consistent with DJ to deny that doxastic justification—or knowledge, as Sylvan (2018) does—entails believing in response to normative reasons one possesses. It is also consistent with DJ to deny that there is a basing requirement on doxastic justification, as Silva (2015a) argues.

²⁶ See, for e.g., (Lord, 2018; Neta, 2019, 2022; Silva, 2015a, 2015b; Sylvan, 2016; Turri, 2010). See also (Carter & Bondy, 2019) for a recent anthology on the basing relation.

²⁷ For more on the category of *normative achievement*, see (Singh, 2020).

Consider again the case of Miss Improper. Miss Improper bases her belief that Mansour is guilty on the sufficient reasons R that she possesses to believe that he is guilty (i.e., (P1 – P4)). Yet, her basing is improper, and BA explains why: Miss Improper’s basing is *not* an exercise of her cognitive abilities, at least not primarily. Unlike Miss Proper, Miss Improper lacks the cognitive ability to infer that Mansour is guilty from (P1 – P4) alone. She needs an intervention from the tea leaves. Because her basing is not an exercise of her cognitive abilities, she is not creditworthy for basing her belief that Mansour is guilty on (P1 – P4), and her basing her belief on (P1 – P4) is not a normative achievement. Since the consequent of BA is false, so is its antecedent: Miss Improper’s basing is improper.

Consider also Sam in El Clasico. Sam believes that Real Madrid lost yesterday. Her belief is based on the sufficient reasons R she possesses to believe so. Yet, her belief is not *properly* based on R. Again, BA explains why: Sam does not competently derive her belief that Real Madrid lost yesterday from the sufficient reasons R that she possesses. Rather, she appeals to an absurd inference rule. So, Sam basing her belief on R is not an exercise of her cognitive abilities. That explains why Sam is not creditworthy for basing her belief that Real Madrid lost yesterday on R. Since she is not creditworthy for doing so, her basing is not a normative achievement. Again, the consequent of BA is false. So, its antecedent is false, too: Sam’s basing is improper.

That having a doxastically justified belief is a normative achievement gets further support from the fact that properly basing one’s belief that *p* on the sufficient reasons R that one possesses to believe that *p* is *but one instance* of correctly responding to sufficient reasons to ϕ

that one possesses, which is more obviously a normative achievement.²⁸ If you properly base your belief that p on the sufficient reasons R that you possess to believe that p , then you are correctly responding to the normative reasons that you possess. So, having a doxastically justified belief is a normative achievement.

We have seen that achievements are successes due to abilities. Here is another platitude: If your success is a matter of luck, it does not constitute an achievement. That is, if your belief that p is doxastically justified, then you believe what you have sufficient reasons to believe, and it is no lucky accident that you so believe. If you get lucky in believing what you have sufficient reasons to believe, you are not creditworthy for so believing, and believing so does not constitute a normative achievement. So, in addition to a success condition on doxastic justification, there is also an anti-luck condition. That is, S 's belief that p is doxastically justified only if:

SUCCESS_{DJ}: S believes what they have sufficient reasons to believe, and

ANTI-LUCK_{DJ}: It is no lucky accident that S believes what they have sufficient reasons to believe.

Consider again the case of Miss Improper. Miss Improper believes what she has sufficient reasons to believe. So, she satisfies SUCCESS_{DJ}. Yet, she is not creditworthy for so believing. As stated earlier, what explains this is the fact that her cognitive success is not due to her cognitive abilities. She needs an intervention from the tea leaves. We are now in a position to say something more. Miss Improper is not creditworthy for believing what she has sufficient reasons to believe because there is a real sense in which Miss Improper was lucky in believing what she has sufficient reasons to believe. There are very close possible worlds in which she

²⁸ What correctly responding to possessed normative reasons amounts to is itself a matter of controversy. For a plausible account, see (Lord, 2018), although, cf. (Kieseewetter, 2017; Neta, 2019). For a fuller defense of the claim that correctly responding to normative reasons that one possesses is a normative achievement, see (Lord, 2018, chs. 3 and 6).

reads off the tea leaves that her evidence supports believing that Mansour is innocent, even though it supports believing that he is guilty. In these close possible worlds, Miss Improper believes what she has insufficient reasons to believe, and not what she has sufficient reasons to believe. So, Miss Improper does not satisfy ANTI-LUCK_{DJ}, and thus her belief is not doxastically justified.

Consider also the case of El Clasico. Sam believes what she has sufficient reasons to believe. Yet, she is not creditworthy for so believing. As stated above, this is because her belief is not due to her cognitive abilities. But we can now say more. It is a matter of luck that Sam believes what she has sufficient reasons to believe. Luckily for Sam, she uses the Explosion rule on the day of El Clasico. But there are very close possible worlds in which Sam uses the rule on the true proposition that Barcelona won yesterday to form her belief that Real Madrid lost, but where she lacks sufficient reasons to believe it.²⁹ So, Sam does not satisfy ANTI-LUCK_{DJ}, and her belief is thus not doxastically justified.

The fact that we can explain why Miss Improper's and Sam's beliefs are not doxastically justified by appealing to either BA, as we did earlier, or ANTI-LUCK_{DJ}, as we did above, is by no means a coincidence. By properly basing one's belief on the sufficient reasons R one possesses to believe that *p*, one sees to it that it is no lucky accident that one believes what one has sufficient reasons to believe. That is, one way to think about the anti-luck condition on doxastic justification is simply in terms of 'proper basing', such that it is no lucky accident that S believes what S has sufficient reasons R to believe if:

PROPER BASING_{DJ} (PB_{DJ}): S's belief that *p* is properly based on the sufficient reasons R that S possesses to believe that *p*.

²⁹ For example, there are very close possible (non-El Clasico) worlds in which Sam truly believes that Barcelona played and beat Valencia yesterday, and Real Madrid played against Las Palmas. (Again, if it helps, we can include this last bit in Sam's epistemic ken.)

Let us take stock. I have argued in this section that having a doxastically justified belief is a normative achievement. It is success (i.e., believing what there is sufficient reasons to believe) due to cognitive abilities (i.e., properly basing one's belief that p). In virtue of being a normative achievement, doxastic justification excludes luck. I have argued that when a subject S has sufficient reasons to believe that p , one way to see to it that believing what there is sufficient reasons to believe is no lucky accident is by properly basing one's belief that p on the sufficient reasons that one possesses to believe that p . So, PB_{DJ} acts as an anti-luck condition on doxastic justification.

3.2. Doxastic justification and modal luck

At this point, one might wonder whether I have got this right. In the case of knowledge, proponents of MVE argue that ACHIEVEMENT cannot adequately capture the anti-luck platitude on knowledge, since ACHIEVEMENT is compatible with knowledge-undermining luck. According to them, adequately capturing the anti-luck platitude requires an additional modal principle like SAFETY. So, one might equally think that in order to properly capture the anti-luck condition on doxastic justification, we need a similar, SAFETY-like condition.

What would such a principle look like? In the case of knowledge, SAFETY requires there to be no very close possible world, and at most few close possible worlds, in which your belief is formed in the same way as in the actual world, but where your belief is false. From the perspective of virtue epistemology, SAFETY requires that one's cognitive success not have easily been a cognitive failure. We have already seen what success amounts to with respect to doxastic justification:

SUCCESS_{DJ}: S believes what they have sufficient reasons to believe.³⁰

With SUCCESS_{DJ} in mind, we can formulate a SAFETY-like principle for doxastic justification thusly:

SAFETY_{DJ}: S's belief that *p* is doxastically justified only if in most nearby possible worlds in which S properly bases their belief that *p* on reasons R, and in all very close nearby possible worlds in which S properly bases their belief that *p* on R, S continues to believe what they have sufficient reasons to believe.

In the remainder of this section, I will show that doxastic justification is subject to modal luck, and thus not constrained by SAFETY_{DJ}. I do this by offering an example of a subject holding a doxastically justified belief that *p*, but where there are very close nearby possible worlds in which S properly bases their belief that *p* on reason R, but where R does not constitute sufficient reasons for S to believe that *p*.

Consider the case of Fahim:

RED WALLS: Today is Fahim's birthday, and Efe is planning a surprise party at her new art gallery, whose walls are currently red. While Fahim has yet to see her new art gallery, Efe knows that Fahim would very much enjoy the walls as they are. Still, for what she (incorrectly) thinks would provide a more saturated look, Efe decides to install red halogen spotlights, which make the walls appear red, hoping that Fahim will be too enthralled by the surprise party to notice. A few minutes before Fahim arrives, however, the bulbs burn out. Upon arrival, Fahim perceives the walls as red, and properly bases his belief that the walls are red on the fact that they appear red.

RED WALLS is a case of what I am calling 'modal luck'. I take it to be undeniable that Fahim's belief that the walls are red is actually properly based on the sufficient reason (the appearance fact) he has to believe that the walls are red. According to DJ, Fahim's belief is doxastically

³⁰ At this point, one might also wonder why success for doxastic justification is not properly basing one's belief that *p* in the sufficient reasons R one possesses to believe that *p*. In short, the answer is that we want to maintain a distinction between *success* and *achievement*. Not every success amounts to an achievement. As we have seen, some successes are due to luck. But when one's success is due to one's abilities, that constitutes an achievement. So, properly basing one's belief that *p* in the sufficient reasons R one possesses to believe that *p* is more than a success. It is an achievement.

justified. So, his cognitive success (i.e., his believing what there is sufficient reason to believe) amounts to a normative achievement.

But Fahim is *lucky* to be in a possible world in which his belief amounts to a normative achievement. There are very close possible worlds in which (i) – the bulbs do not burn out, (ii) – Fahim properly bases his belief that the walls are red on the fact that they appear red, and (iii) – Fahim is in a position to know that the walls are bathed in red light.³¹ Yet, in these very close possible worlds, Fahim’s reason for his belief—i.e., the fact that the walls appear red—is not a sufficient reason to believe that the walls are red. Given that Fahim is in a position to know that the walls are bathed in red light, he has higher-order evidence that his first-order evidence is undermined. Since Fahim’s evidence is undermined in this way, he believes what he has insufficient reason to believe. Since there are very close possible worlds in which Fahim believes what he has insufficient reasons to believe, his actual success could have easily been a cognitive failure.

I take this to be clear evidence that doxastic justification, understood as a cognitive normative achievement, is not susceptible to SAFETY-like principles like SAFETY_{DJ}. In the following section, I show that performing a morally worthy action—another kind of normative achievement—is also susceptible to modal luck and thus unconstrained by a similar SAFETY-like principle.

3.3. Moral worth

³¹ I add this because, like Lycan (1977), I am skeptical that things outside of one’s epistemic ken can act as undercutting defeaters. I am sympathetic to the idea that one being in a position to know that p is sufficient for p to be in one’s epistemic ken, but nothing here hinges on any specific account of epistemic possession. (See (Lord, 2018, chs. 3 & 4) on reason possession.) If that the walls are bathed in red light is in Fahim’s epistemic ken, then there is little doubt that Fahim’s justification is undercut.

An agent's morally right action has moral worth if and only if it is done for the right reasons.

There is a live debate about what constitutes the right reasons. According to the right-making features view (RMFV), the right reasons are the reasons that make the action right.³² According to rightness-itself view (RIV), the right reason is simply the fact that the action is morally right.³³

While there is considerable disagreement between the views, they are both motivated in large part by a desire to come up with an account of acting for the right reasons (whatever those turn out to be), that can eliminate the possibility of an agent accidentally doing the right thing. The reason for this should now be obvious: Performing a morally worthy action is a normative achievement.³⁴ If your action has moral worth, then you are creditworthy for doing the right thing. But if you only do the right thing accidentally, you are not creditworthy for doing the right thing.³⁵ And if you are not creditworthy for doing the right thing, you doing the right thing is not a normative achievement. So, as with doxastic justification, there is both a success condition and an anti-luck condition on moral worth, such that an agent A performs a morally worthy action only if:

SUCCESS_{MW}: A does what is morally right,³⁶ and

ANTI-LUCK_{MW}: It is no lucky accident that A does what is morally right.

³² For proponents of RMFV, see, e.g., (Arpaly, 2003; Markovits, 2010).

³³ For proponents of RIV, see, e.g., (Johnson King, 2020; Sliwa, 2016)

³⁴ This I take to be uncontroversial. In any case, I also leave this for the most part as an undefended assumption.

³⁵ This, of course, is an insight that goes at least back to Kant (2018). In the Preface to the *Groundwork*, Kant writes: "In the case of what is to be morally good, it is not enough that it *conform* to the moral law, but it must also happen *for the sake of this law*; otherwise, that conformity is only contingent and precarious, because the unmoral ground will now and then produce lawful actions, but more often actions contrary to the law (Ak 4: 390).

³⁶ That success for moral worth involves doing the right thing makes it clear that performing a morally worthy action is not *merely* a cognitive achievement. It is also a practical achievement, and practical achievements depend on things outside of one's cognition. Still, arguably, in order to do the right thing for the right reason, you need to *intend* to do the right thing, and unless the world transpires against you, your action will have moral worth if you intend to do the right thing for the right reason, which is surely a cognitive (and normative) achievement.

As stated above, proponents of RMFV and RIV alike agree that in order to guarantee that it is not a matter of luck that an agent does what is right, the agent needs to do the right thing *for* the right reasons (whatever those turn out to be). Here again, I wish to remain as ecumenical as possible, both with respect to which reasons are the right ones, but also with respect to what it is to act for the right reason. Still, I take it that proponents of either view are amenable to understanding the anti-luck condition on moral worth in terms of ‘proper basing’, such that A doing the right thing has moral worth only if:

PROPER BASING_{MW} (PB_{MW}): A’s doing what is morally right is properly based on the right reasons.

3.4. Moral worth and modal luck

Again, someone might object that PB_{MW} on its own cannot act as the anti-luck condition on moral worth. One might worry that, just as ACHIEVEMENT is taken to be compatible with knowledge-undermining luck, PB_{MW} is compatible with moral worth-undermining luck. Like knowledge, some might think that capturing the anti-luck condition on moral worth is going to require a modal principle like SAFETY.

What would such a SAFETY-like principle look like? Again, from the perspective of virtue epistemology, SAFETY requires that one’s cognitive success not have easily been a cognitive failure. So, we should expect a SAFETY-like principle for moral worth to require that one’s normative success not have easily been a normative failure. We have already seen above what constitutes normative success for moral worth:

SUCCESS_{MW}: A does what is morally right.

With SUCCESS_{MW} in mind, we can formulate a SAFETY-like principle for moral worth:

SAFETY_{MW}: A’s ϕ -ing has moral worth only if in most nearby possible worlds in which A properly bases their ϕ -ing on the right reason R, and in all very close

nearby possible worlds in which A properly bases their ϕ -ing on R, A's ϕ -ing is morally right.

In the remainder of this section, I will show that like doxastic justification, moral worth is subject to modal luck, and thus not constrained by SAFETY_{MW}. I do this by offering an example of an agent A performing a morally worthy action, but where there are very close nearby possible worlds in which A properly bases their ϕ -ing on reason R, but where ϕ -ing is morally wrong.

Consider the case of Danny.

DOG-LOVER: Danny is a fanatic dog-lover who works for the Animal Humane Society. She lives in a duplex house, and next door is an elderly couple who rarely leaves the house. Danny is currently fostering Cinnamon, an 11-year-old Shepherd-mix. Danny drives home every day during her lunch hour to walk Cinnamon. But today is an unusual day: the elderly couple decides to go to the one-day bazaar at the local church, letting Danny know as they leave in the morning that they'll be out for the day. On her way home during her lunch hour, Danny sees that the duplex is engulfed in smoke. Danny rushes inside, at some personal risk, and rescues Cinnamon from the smoke-filled building.³⁷

Suppose that rushing inside the building to rescue Cinnamon is the right thing for Danny to do, and that Danny's action is properly based on the right reasons R (either on the facts about the features that make Danny's action morally right, or simply on the fact her action is morally right). Under these assumptions, it strikes me as obvious that Danny's action has moral worth. Danny is, after all, creditworthy for doing the right thing. Her action is a normative achievement.

But now consider some important background information: Since Danny is a fanatic dog-lover, she would have rushed in to save Cinnamon for the same reasons R, instead of the elderly couple, had the elderly couple not gone to the bazaar and stayed home. And let us stipulate that,

³⁷ This example is largely inspired by an example from Markovits (2010, p. 210), where a fanatic dog-lover performs a dangerous rescue operation to save the lives of strangers, but would not have done so had their dog required their heroics at the same time. Arpaly (2003), who defends a quality-of-will account of moral worth, takes examples like DOG-LOVER to show that agents like Danny lack moral worth, at least to some degree. Markovits's point with the example—a point with which I agree—is to press the importance of distinguishing between the moral worth of an action and an agent's character. See also (Johnson King, 2020) who presses the same important point.

had the elderly couple stayed home, rushing inside the building to save Cinnamon instead of the elderly couple would not have been the morally right thing to do.

DOG-LOVER is also a case of modal luck. Danny is *lucky* to be in a possible world in which her action amounts to normative achievement. There are very close possible worlds in which (i) – the elderly couple stays home, (ii) – Danny properly bases her action on R, and (iii) – Danny is in a position to know that the elderly couple is home. Yet, in these very close possible worlds, Danny’s action—i.e., rushing in to save Cinnamon rather than the elderly couple—is not morally right. Since there are very close possible worlds in which Danny’s action is not morally right, this means that her actual normative success could have easily been a normative failure.

I take DOG-LOVER to show that moral worth is not constrained by a SAFETY-like principle like SAFETY_{MW}. But this, I must admit, might not convince proponents of RIV. If the right reason R involved in moral worth is the fact that the action is right, as proponents of RIV argue, then Danny’s reason for action in rushing in to save Cinnamon in the actual world differs from her reason for action in those very close possible worlds in which the elderly couple stays home. If RIV is true, then in the actual world, Danny’s reason for action is the fact that rushing in to save Cinnamon is the right thing to do. But that cannot be Danny’s reason for action in the close possible worlds in which the elderly couple stays home since, *ex hypothesi*, it is not a fact that saving Cinnamon in those possible worlds is morally right. So, DOG-LOVER is not a counterexample to SAFETY_{MW} after all.

I do not have the space here to offer a full response to this objection. One thing to note is that the objection rests on the assumption that motivating reasons are factive, and it is controversial whether motivating reasons are factive.³⁸ If motivating reasons are factive and RIV

³⁸ For arguments against the factivity of motivating reasons (either for belief or action), see, e.g., (Comesaña & McGrath, 2014; Dancy, 2000; Fantl, 2015; Schroeder, 2008; Singh, 2019).

is true, then DOG-LOVER is not a counterexample to SAFETY_{MW}. But it is also controversial whether RIV is true. If RMFV is true instead, then regardless of whether motivating reasons are factive, DOG-LOVER is a counterexample to SAFETY_{MW}. And even if RIV is true, such that properly basing one's right action on facts about the features that make the action right lacks moral worth, it is nevertheless a normative achievement. It is, after all, an instance of correctly responding to the moral reasons one possesses, which is a normative achievement, and that is really all I need. So, here is yet another kind of normative achievement (which may or may not constitute moral worth) that is susceptible to modal luck and thus not constrained by a SAFETY-like principle.

4. Knowledge and modal luck

I have argued that two important normative achievements that are rightly taken to exclude some kinds of luck are nevertheless susceptible to modal luck, and are thus unconstrained by SAFETY-like principles. So, if knowledge is a normative achievement, then, short of a principled reason why we should expect knowledge to be different in this regard, we should not be surprised if knowledge is also susceptible to modal luck.

Is knowledge a normative achievement? Virtue epistemologists agree that knowledge is, at least in part, a cognitive achievement. Is it also a *normative* achievement? Arguably, some instances of knowledge are normative achievements. At the very least, some instances of knowledge *entail* normative achievements: many cases of knowledge are cases of one properly basing one's true belief on the sufficient reasons R one possesses to believe that *p*, which—as we have seen—are cases of normative achievements. Consider again the case of the bisque. I know that eating the bisque will make me ill. So, I have a true belief that the bisque will make me ill,

and that true belief is properly based on the sufficient reasons I possess to so believe. Since my true belief is properly based, it is a normative achievement. So, at least some instances of knowledge are, or entail, normative achievements.

That some instances of knowledge are normative achievements gets further support from the fact that normative achievements are creditworthy, and knowledge is something for which one is creditworthy. From the perspective of virtue epistemology, knowledge is cognitive success due to cognitive abilities. When a cognitive success is due to your cognitive abilities, you are creditworthy for that success. And when the ability involves properly basing one's true belief that p on the sufficient reasons one possesses to believe that p , that constitutes a normative achievement. Hence, short of a principled, non-*ad-hoc* reason to expect knowledge to be different in this regard, we should expect knowledge (at least some instance) to also be susceptible to modal luck.

And when we think about certain cases, that is in effect what we see. Consider a variant of RED WALLS:

RED WALLS*: Today is Fahim's birthday, and Efe is planning a surprise party at her new art gallery, whose walls are currently red. Efe knows that Fahim would very much enjoy the walls as they are, but a new installation is scheduled to go up tomorrow, so Efe wants to put up a fresh coat of white paint on the walls before the party guests arrive this evening. Still, Efe wants to make tonight's party special for Fahim. So, in addition to buying (what she thinks is) a gallon of white paint at the hardware store, Efe buys a few red spotlights which, when turned on, make white walls appear red. Back at the gallery, Efe installs the lights before painting the walls. As she opens the gallon of paint, Efe realizes that the hardware store clerk accidentally sold her a gallon of Prussian blue. Since there is not enough time to exchange the gallon, paint the room white, and have it dried before the first guests arrive, Efe decides to forgo using the spotlights, apprehensively leaving the painting for tomorrow. Upon arrival, Fahim perceives the walls as red, and bases his belief that the walls are red on the fact that the walls appear red.³⁹

³⁹ Notice that, like the other cases of modal luck, the strategy deployed in RED WALLS* involves putting spoilers in very close possible worlds, but not in the actual world. This is precisely the strategy deployed in what I take to be successful examples of unsafe knowledge. See, e.g., (Comesaña, 2005; Kelp, 2009; McBride, 2011).

RED WALLS* is another case of modal luck. I take it to be obvious that Fahim knows that the walls are red. Yet, Fahim is *lucky* to be in a possible world in which his true belief amounts to knowledge. There are very close nearby possible worlds in which Fahim believes that the walls are red, where his belief is properly based on the appearance-fact, but where his belief is false. These are the worlds in which the hardware store clerk does not make a mistake in selling Efe a gallon of Prussian blue. If there is knowledge in RED WALLS*, then SAFETY is not a condition for knowledge. If SAFETY is not a condition for knowledge, then it cannot be used to adequately capture the anti-luck platitude that knowledge excludes luck. Virtue epistemologists must look elsewhere. My suggestion is to stick with ACHIEVEMENT.

5. Concluding remarks: Environmental vs. modal luck

It is important to remember that while I think that there are cases of unsafe knowledge—cases like RED WALLS*—my argument did not essentially rely on such examples. Rather, my argument proceeded by way of analogy. We looked at paradigmatic cases of normative achievements, and we saw that, while these normative achievements are incompatible with (some kinds of) luck, they are nevertheless susceptible to modal luck. In virtue of this, they are not constrained by SAFETY-like modal principles. And since some instances of knowledge either are, or entail, normative achievements, the arguments from section 3 provided us with good reason to expect knowledge to also be susceptible to modal luck, and thus unconstrained by SAFETY. That we can cook up some convincing cases of unsafe knowledge is simply proof of concept.

If modal luck is not knowledge-undermining, as I have argued above, where does that leave the knowledge-undermining status of environment luck? I do not have the space here to

offer a satisfactory answer, but there are at least two available options. First, one might think that there is no substantive difference between modal and environmental luck. After all, in the actual world, Henry is looking at the only real barn in the county. So, Henry is *lucky* to be in a possible world in which he is looking at a real barn: there are very close nearby worlds in which Henry is looking at a barn façade. If there is no substantive difference between modal and environmental luck, then the argument above provides reason to deny that environmental luck is knowledge-undermining. That would not be so bad, in my view.

On the other hand, there does seem to be a difference between the two. In FAKE BARNs, it is natural to think that what spoils things for Henry—i.e., what some think undermines knowledge in cases of environmental luck—is in the actual world. Actually, Henry is driving through Fake Barn County. That is what is supposed to spoil things for Henry. But in cases of modal luck, the spoiler is not in the actual world. It is only in very close nearby possible worlds. That might be what explains a willingness to grant knowledge in cases of modal luck. If this appearance turns out to not be illusory, this could go some way to explain why environmental luck, but not modal luck, is knowledge-undermining. In any case, given that we have good reason to think that knowledge is compatible with modal luck, we have good reason to reject SAFETY as a condition for knowledge, one that can adequately capture the anti-luck platitude on knowledge.⁴⁰

References

Arpaly, N. (2003). *Unprincipled virtue: An inquiry into moral agency*. Oxford University Press.

⁴⁰ I would like to thank audience members at the 2023 Zürich-Uppsala Meeting on Epistemic Normativity for some valuable feedback on an earlier version of this paper, and to Simon-Pierre Chevarie-Cossette for some helpful discussion on how to frame some of the ideas in this paper.

- Carter, J. A. (2013). A Problem for Pritchard's Anti-Luck Virtue Epistemology. *Erkenntnis*, 78(2), 253–275. <https://doi.org/10.1007/s10670-011-9315-x>
- Carter, J. A., & Bondy, P. (Eds.). (2019). *Well-founded belief: New essays on the epistemic basing relation*. Routledge.
- Chisholm, R. M. (1977). *Theory of knowledge* (2d ed.). Englewood Cliffs, N.J. : Prentice-Hall.
- Colaço, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic Intuitions in Fake-Barn Thought Experiments. *Episteme*, 11(2), 199–212. <https://doi.org/10.1017/epi.2014.7>
- Comesaña, J. (2005). Unsafe Knowledge. *Synthese*, 146(3), 395–404. <https://doi.org/10.1007/s11229-004-6213-7>
- Comesaña, J., & McGrath, M. (2014). Having False Reasons. In C. Littlejohn & J. Turri (Eds.), *Epistemic Norms* (pp. 59–79). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199660025.003.0004>
- Dancy, J. (2000). *Practical reality*. Oxford.
- Fantl, J. (2015). What Is It to Be Happy That P? *Ergo, an Open Access Journal of Philosophy*, 2(20201214). <https://doi.org/10.3998/ergo.12405314.0002.012>
- Firth, R. (1978). Are Epistemic Concepts Reducible to Ethical Concepts? In A. Goldman & J. Kim (Eds.), *Values and morals: Essays in honor of William Frankena, Charles Stevenson, and Richard Brandt*. Dordrecht.
- Gendler, T. S., & Hawthorne, J. (2005). The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124, 331–352.
- Gettier, E. L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–123. <https://doi.org/10.1093/analys/23.6.121>

- Goldman, A. (1976). Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73(20), 771. <https://doi.org/10.2307/2025679>
- Goldman, A. (1992). *Liaisons: Philosophy meets the cognitive and social sciences*. MIT Press.
- Goldman, A. (2009). Williamson on Knowledge and Evidence. In P. Greenough & D. Pritchard (Eds.), *Williamson on Knowledge*. Oxford University Press.
- Goldman, A. (2012). *Reliabilism and Contemporary Epistemology: Essays*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199812875.001.0001>
- Greco, J. (2003). Knowledge as Credit for True Belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (pp. 111–134). Oxford University Press.
- Greco, J. (2009). Knowledge and success from ability. *Philosophical Studies*, 142(1), 17–26. <https://doi.org/10.1007/s11098-008-9307-0>
- Greco, J. (2012). A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research*, 85(1), 1–26. <https://doi.org/10.1111/j.1933-1592.2011.00567.x>
- Jarvis, B. (2013). Knowledge, Cognitive Achievement, and Environmental Luck. *Pacific Philosophical Quarterly*, 94(4), 529–551. <https://doi.org/10.1111/papq.12012>
- Jarvis, B., & Ichikawa, J. J. (2013). *The rules of thought* (1st edition.). Oxford : Oxford University Press.
- Johnson King, Z. A. (2020). Accidentally Doing the Right Thing. *Philosophy and Phenomenological Research*, 100(1), 186–206. <https://doi.org/10.1111/phpr.12535>
- Kallestrup, J., & Pritchard, D. (2013). Robust Virtue Epistemology and Epistemic Dependence. In T. Henning & D. P. Schweikard (Eds.), *Knowledge, Virtue, and Action: Putting Epistemic Virtues to Work* (pp. 209–226). Routledge.

- Kallestrup, J., & Pritchard, D. (2014). Virtue Epistemology and Epistemic Twin Earth. *European Journal of Philosophy*, 22(3), 335–357. <https://doi.org/10.1111/j.1468-0378.2011.00495.x>
- Kallestrup, J., & Pritchard, D. (2016). Dispositional Robust Virtue Epistemology versus Anti-luck Virtue Epistemology. In M. Á. Fernández Vargas (Ed.), *Performance Epistemology* (pp. 31–50). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198746942.003.0003>
- Kant, I. (2018). *Groundwork for the Metaphysics of Morals* (A. W. Wood, Trans.). United States: Yale University Press.
- Kelp, C. (2009). Knowledge and Safety. *Journal of Philosophical Research*, 34, 21–31. https://doi.org/10.5840/jpr_2009_1
- Kelp, C. (2013). Knowledge: The Safe-Apt View. *Australasian Journal of Philosophy*, 91(2), 265–278. <https://doi.org/10.1080/00048402.2012.673726>
- Kiesewetter, B. (2017). *The normativity of rationality* (First edition). Oxford University Press.
- Krasser, H. (1995). Dharmottara’s Theory of Knowledge in his “Laghuprāmāṇyaparīkṣā.” *Journal of Indian Philosophy*, 23(2), 247–271.
- Kvanvig, J. L. (2003). Propositionalism and the Perspectival Character of Justification. *American Philosophical Quarterly (Oxford)*, 40(1), 3–17.
- Littlejohn, C. (2014). Fake barns and false dilemmas. *Episteme*, 11(4), 369–389. <https://doi.org/10.1017/epi.2014.24>
- Lord, E. (2018). *The Importance of Being Rational*. Oxford University Press. <https://doi.org/10.1093/oso/9780198815099.001.0001>

- Lord, E., & Sylvan, K. (2019). Prime Time (for the Basing Relation). In J. A. Carter & P. Bondy (Eds.), *Well Founded Belief: New Essays on the Epistemic Basing Relation* (pp. 141–173). Routledge.
- Lycan, W. G. (1977). Evidence one does not possess. *Australasian Journal of Philosophy*, 55(2), 114–126. <https://doi.org/10.1080/00048407712341141>
- Lycan, W. G. (2006). *On the Gettier Problem problem*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780199273317.003.0009>
- Markovits, J. (2010). Acting for the Right Reasons. *The Philosophical Review*, 119(2), 201–242. <https://doi.org/10.1215/00318108-2009-037>
- McBride, M. (2011). Reply to Pardo: Unsafe Legal Knowledge? *Legal Theory*, 17(1), 67–73. <https://doi.org/10.1017/S1352325211000048>
- Neta, R. (2019). The Basing Relation. *The Philosophical Review*, 128(2), 179–217. <https://doi.org/10.1215/00318108-7374945>
- Neta, R. (2022). Justification Ex Ante and Ex Post: Why We Need Both Notions, and Neither is Reducible to the Other. In P. Silva & L. R. G. Oliveira (Eds.), *Propositional and Doxastic Justification* (1st ed., pp. 59–75). Routledge. <https://doi.org/10.4324/9781003008101-6>
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed). Rowman & Littlefield Publishers.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158(3), 277–297. <https://doi.org/10.1007/s11229-006-9039-7>

- Pritchard, D. (2009). Safety-Based Epistemology: Whither Now? *Journal of Philosophical Research*, 34, 33–45. https://doi.org/10.5840/jpr_2009_2
- Pritchard, D. (2012a). Anti-Luck Virtue Epistemology. *Journal of Philosophy*, 109(3), 247–279. <https://doi.org/10.5840/jphil201210939>
- Pritchard, D. (2012b). In defence of modest anti-luck epistemology. In K. Becker & T. Black (Eds.), *The Sensitivity Principle in Epistemology* (pp. 173–192). Cambridge University Press. <https://doi.org/10.1017/CBO9780511783630.014>
- Pritchard, D. (2014). The Modal Account of Luck. *Metaphilosophy*, 45(4–5), 594–619. <https://doi.org/10.1111/meta.12103>
- Pritchard, D. (2015). Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172(1), 93–111. <https://doi.org/10.1007/s11098-014-0374-0>
- Pritchard, D. (2016). Epistemic risk. *The Journal of Philosophy*, 113(11), 550–571.
- Radford, C. (1966). Knowledge: By Examples. *Analysis*, 27(1), 1. <https://doi.org/10.2307/3326979>
- Russell, B. (2013). *Human Knowledge: Its Scope and Value* (0 ed.). Routledge. <https://doi.org/10.4324/9780203713952>
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1), 57–71. <https://doi.org/10.1007/s11098-007-9102-3>
- Silva, P. (2015a). Does Doxastic Justification Have a Basing Requirement? *Australasian Journal of Philosophy*, 93(2), 371–387. <https://doi.org/10.1080/00048402.2014.962553>
- Silva, P. (2015b). On Doxastic Justification and Properly Basing One's Beliefs. *Erkenntnis*, 80(5), 945–955. <https://doi.org/10.1007/s10670-014-9690-1>

- Silva, P. (2023). *Awareness and the Substructure of Knowledge* (1st ed.). Oxford University PressOxford. <https://doi.org/10.1093/oso/9780192869548.001.0001>
- Silva, P., & Oliveira, L. R. G. (2023). Propositional Justification and Doxastic Justification. In M. Lasonen-Aarnio & C. Littlejohn (Eds.), *The Routledge Handbook of the Philosophy of Evidence* (1st ed., pp. 395–407). Routledge. <https://doi.org/10.4324/9781315672687-36>
- Singh, K. (2019). Acting and Believing Under the Guise of Normative Reasons. *Philosophy and Phenomenological Research*, 99(2), 409–430. <https://doi.org/10.1111/phpr.12497>
- Singh, K. (2020). *Rational Agency and Normative Achievement*. University of North Carolina at Chapel Hill.
- Sliwa, P. (2016). Moral Worth and Moral Knowledge. *Philosophy and Phenomenological Research*, 93(2), 393–418. <https://doi.org/10.1111/phpr.12195>
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199297023.001.0001>
- Swain, M. (1979). Justification and the Basis of Belief. In G. S. Pappas (Ed.), *Justification and Knowledge: New Studies in Epistemology* (pp. 25–49). Springer Netherlands. https://doi.org/10.1007/978-94-009-9493-5_2
- Sylvan, K. (2016). Epistemic Reasons II: Basing. *Philosophy Compass*, 11(7), 377–389. <https://doi.org/10.1111/phc3.12329>
- Sylvan, K. (2018). Knowledge as a Non-Normative Relation. *Philosophy and Phenomenological Research*, 97(1), 190–222. <https://doi.org/10.1111/phpr.12371>

Turri, J. (2010). On the Relationship between Propositional and Doxastic Justification.

Philosophy and Phenomenological Research, 80(2), 312–326.

<https://doi.org/10.1111/j.1933-1592.2010.00331.x>

Turri, J. (2011). Manifest Failure: The Gettier Problem Solved. *Philosophers' Imprint*, 11, 1–

11.

Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical*

foundations of knowledge. Cambridge University Press.