

Draft. Forthcoming, *Episteme*.

The Epistemic Analysis of Luck¹

Gregory Stoutenburg

Abstract:

Duncan Pritchard has argued that luck is fundamentally a modal notion: an event is lucky when it occurs in the actual world, but not more than half of the relevant nearby possible worlds. Jennifer Lackey has provided a powerful counterexample to Pritchard's account. Neil Levy has responded to Lackey by offering a modal account of luck which attempts to respect the intuition that some lucky events occur in more than half of the relevant nearby possible worlds. But his account rejects that events which are as likely as those in Lackey's examples are lucky. Instead, they are merely fortunate. I argue that Levy's argument to this effect fails. I then offer a substitute account of the improbability condition which respects this intuition. This condition says that the relevant notion of probability for luck-attribution is epistemic, not modal.

Introduction

An increasingly pressing problem in epistemological, ethical, and metaphysical circles is that important issues appear to hinge on the concept of luck, for which we lack a satisfactory analysis. Knowledge is widely thought to be incompatible with beliefs that are only luckily true; praiseworthy or blameworthy moral conduct is thought to depend on the absence of moral luck;

¹ Thanks to Clinton Packman, Ryan Cobb, Samuel Taylor, Chris Dyer, Heather Stoutenburg, Bryan Appley, Ali Hasan, Brady Hoback, Patrick Beach, Mylan Engel, an anonymous referee for *Episteme*, and audiences at the Midsouth Philosophy Conference 2014 and a University of Iowa Graduate Philosophical Society Salon for comments and helpful discussion. I especially thank Brian Collins and Richard Fumerton for detailed comments on multiple drafts of this paper.

the possibility of free action is believed to hinge on the absence of luck in such action. Serious attempts to analyze the concept of luck are beginning to appear. The most prominent accounts analyze luck in terms of the proportion of possible worlds in which some event occurs. In this paper I defend a novel analysis of luck. I argue that lucky events are those which are significant for a subject and epistemically uncertain. I proceed as follows. I briefly identify the key intuitions which drive our practice of luck attribution. I then show how the modal account tries to capture these intuitions but fails to do so, and how my proposed epistemic analysis of luck adequately captures our key intuitions about luck.

Luck and Luck Attribution

This paper concerns the conditions under which an event (or, indirectly, a person) is lucky. On the assumption that luck attributions are often true, we will be able to motivate a correct analysis of luck by means of searching for the conditions under which we are inclined to attribute luck. It is not very difficult to get someone to agree that many commonplace events are at least a little bit lucky. Our talk indicates that luck comes in degrees. In ordinary language we say some event was “very lucky” or “a little lucky”. These and similar expressions are used to approximate the degree to which some event had the property of being lucky. When some event is even a little bit lucky, it is *true* to say it is lucky: precisely, it is a little bit lucky. The truth of a luck-attribution depends only on whether the event referred to in the proposition expressed by the speaker has the property of being lucky. But the *appropriateness* of a luck-attribution has other conditions. For example, someone might be just the slightest bit lucky to score a goal, but on account of the person’s skill, her goal is unsurprising. In that context, the goal has the property of being lucky, and it would be true to attribute luck to her goal, but it would be inappropriate to do so verbally. We probably

have a threshold in mind beyond which we will not call an event “lucky”. Still, her shot is in fact a little bit lucky. Just as it is not always appropriate to assert something when it is true, it is not always appropriate to call something lucky when it is, nor to say it is not when it is not. In this paper, I am primarily concerned with the property of luck, and so with the truth of luck-attributions. I will only occasionally discuss the issue of appropriateness.²

The Intuitions of Luck

The concept of luck is an ordinary concept. People freely talk about events being lucky. There are a few intuitions that guide our practice of luck-attribution; I attempt to identify them in this section. It is easy to think of examples of luck. Here is one: you are walking along a boardwalk when you stumble upon a crisp fifty-dollar bill. Being an honest and helpful person, you look around to see who might have dropped the bill. There is no one in sight. You wait to see if anyone comes by for the money. No one does. The money is yours. Most will agree that you got lucky when you found the money. You had no reason to think you would stumble across the money when you went for a stroll. The notion that lucky events are surprising, unpredictable, and improbable seems to play a key role in our attributions of luck. Finding the bill is also personally significant to you. Perhaps it wouldn’t be significant to Warren Buffett on account of his wealth; how could a few more dollars matter to him?

Another example: you enter a game show which offers contestants multiple-choice questions of increasing difficulty on different subjects. The contestant who gets the most answers

² The distinction between the appropriateness of luck-attributions as contrasted with the truth of luck-attributions is the same as the distinction between truth-conditional properties of utterances and non-truth conditional properties of utterances which are important for other purposes to which we put language. Cf. H.P. Grice (1975).

right wins a dream vacation. When you arrive for the filming, you learn that the show's topic of the day is ornithology. You can barely tell the difference between a seagull and a sparrow. Yet you guess your way through each question and win the prize. Again, you were lucky: your guesses being correct was improbable by anyone's lights, including yours, and your win is great news for you.

When we talk about luck, we talk about events that are lucky *for* someone.³ There is no such thing as an intrinsically lucky event. As we just saw in our two examples, lucky events are also improbable (in some sense) and significant to a subject (Rescher 1995, Pritchard 2005, Levy 2011, Coffman 2007, Steglich-Petersen 2010). The key intuitions driving our luck attributions are improbability and significance. This much is nothing new. Analyses of luck on offer in the literature attempt to capture these intuitions as two individually necessary and jointly sufficient conditions:

(1) An *improbability* condition. Intuitively, lucky events are ones that we think are improbable. When someone says E was lucky, the subject intends at least to convey that E was somehow unlikely.

(2) A *significance* condition. Lucky events are ones that *matter* to the subject of the luck-attribution.

Precisely what view of luck one has depends on how exactly the improbability and significance conditions are understood. One could, in principle, conjoin any interpretation of the improbability condition that one likes with any significance condition one likes. For example, one could take a modal improbability condition and add a subjectivist view of significance which says that some event is significant if and only if a subject actually cares about it. A different view results if the

³ As Coffman (2007) indicates, groups may also be the subjects of luck. In such cases, their luck reduces to the luck of the group's members.

subjectivist condition is understood counterfactually, so that an event is significant if the subject would, with appropriate information and reflection, care about the event. Or the significance condition could be understood objectively, so that what is significant for a subject depends upon preference-independent facts about the subject and her place in the world (Coffman 2007). Similar variation is possible with the improbability condition, as we shall see shortly. Most of the debate about luck in the literature concerns the precise formulation of the modal improbability condition (Pritchard 2005, Coffman 2007, Levy 2012). I now turn to that issue.

Improbability—The Modal Accounts

The intuition behind the modal view of luck is straightforward: lucky events are those which could have easily been different. Lucky events are improbable in that sense. The basic idea is that an event is lucky when it occurs in the actual world but not in the majority of nearby worlds. That account of improbability, hooked up with a significance condition, delivers a modal account of luck. I now briefly introduce Pritchard's account.

Pritchard

One currently influential modal account of the improbability condition is offered in Duncan Pritchard's book *Epistemic Luck*. Here is Pritchard:

L1 If an event is lucky, then it is an event that occurs in the actual world but which does not occur in a *wide class* of the nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world (2005).

For Pritchard, what is improbable about a lucky event is that it occurs *here* but not *there*, for a “wide class” of alternative *theres*. Pritchard is not firm about how wide the wide class has to be, granting that a lucky event may obtain in up to half of the nearest possible worlds.

Jennifer Lackey (2008) has presented a powerful objection to L1, of which the following is a paraphrase.

BURIED TREASURE. Sophie, soon to die, buries a chest of valuables on the northwest corner of an island, a place of deep importance to her, and a place where she hoped roses would sprout in the future. Some time later, Vincent comes to the island to plant a rosebush in his mother's memory, and finds the only suitable location: the place directly above Sophie's buried treasure. He begins digging and finds it.

Here is a case where Vincent is intuitively lucky to find the treasure. He found, unaided, a treasure chest. But, as Lackey has pointed out, the case has been crafted so that Vincent finds the treasure in at least the majority of worlds where the relevant initial conditions to the event are the same: the Vincent and Sophie are acting out of deep-seated dispositions, and it is these dispositions that conspire to help Vincent find the treasure.⁴

BURIED TREASURE shows that Pritchard's L1 is not a necessary condition on luck because lucky events may occur in more than half of the relevant nearby possible worlds. Neil Levy defends a modal account and contests the moral drawn from this case.

⁴ It would be easy to imagine that the dispositions in question are such that a description of the nearby neighborhood of possible worlds is such that Vincent was almost guaranteed to find the treasure. To be sympathetic to Lackey's case, I presume we should not think of the event as quite so likely to happen. The intuitive idea behind the case is just that there can be events which are lucky because they are significant to a subject, but they are likely enough to happen that a modal improbability condition will struggle to allow the intuitive verdict on the case.

Of course, there is a dialectical risk in offering any thought experiment: some people won't be moved. To those unmoved, I have two remarks. First, it seems to me that an effective way to check for the intuitive pull to describe a state of affairs presented in a thought experiment using some important term is to offer a minimal description of the scenario and to see how strong the intuition is that the term applies to that state of affairs. If it is strong, that is at least good evidence that the state of affairs is properly described by the term. So, the fact that we are inclined to say that a man finding buried treasure is lucky (*extremely* lucky!) *prior to* learning of the background is at least evidence that the man finding treasure is, in fact, lucky.

Second, the argument of the paper is ultimately conditional in structure: *if* there are lucky-but-probable events, then the epistemic analysis of luck argued for in this paper is the best hope of accounting for them. Clearly, I think the antecedent of the conditional is true.

I owe thanks to an anonymous referee for encouraging me to say more about this case and its role in the structure of my argument.

Levy

Levy defends a modal account as a condition of “chancy” luck (in contrast with “nonchancy” luck, of which constitutive luck is a species) in his *Hard Luck* (2011).

Event E is chancy if it occurs in the actual world at t_1 , but it fails to occur in a *large enough* proportion of possible worlds obtainable by making no more than a small change to the actual world at t_0 ; and the agent lacks direct control over E’s occurrence (Levy 2011).

Levy makes two modifications to L1. First, he adds a lack of control condition. Second, he argues that whether an event is chancy depends in part upon how significant the event would be for the subject. Levy’s threshold for how improbable an event must be to count as lucky is tied to how significant the event is.⁵ The more significant the event is, the more probable the event is allowed to be while still counting as not happening in a “large enough” percentage of possible worlds. This permits Levy to identify (modally) probable events as lucky.

Levy’s view ably handles the Russian Roulette case (Rescher 1995).

RUSSIAN ROULETTE: Samuel plays and wins one standard round of Russian Roulette: a single trigger pull of a 6-round revolver with only one bullet in the chamber.

Levy acknowledges that anyone who plays and “wins” Russian Roulette is lucky. But the odds of winning are in the player’s favor: there is only a 1/6 chance that the outcome is death and a 5/6 chance that the player finishes the game unscathed. Levy argues that what makes a game of Russian Roulette chancy is that the 5/6 odds of survival are *not high enough* given that the significance of the 1/6 event is very high.

⁵ No machinery is provided to explain what exactly the connection is or how the probability threshold is determined by the significance of the event. Levy claims, “Luck is a vague concept, and an adequate account of it must be vague (in the same way and in the same places) as well” (2011, pp.18)

Although Levy wants to allow that sometimes probable events are lucky, he wants to prevent *highly* likely events from counting as lucky, as he recognizes that allowing the significance condition to swamp *any* probability, no matter the degree, threatens the modal account itself. That motivates his response to BURIED TREASURE: because Vincent would find the treasure in too many possible worlds, he is not lucky. Levy argues that we should accept an error theory regarding this mistake. He argues that we often confuse luck with fortune. Call this the ‘confusion with fortune’ error theory. Fortune, he says, is had when one is the positive or negative beneficiary of some event, even if the event was highly probable.⁶ Fortunate events are non-lucky events with luck in their proximate causes. In BURIED TREASURE, Levy argues, Vincent is not lucky to find the treasure, he is merely *fortunate* to do so.

I agree with Lackey that finding buried treasure is “a paradigmatic instance” of a lucky event (2008). Philosophers who agree that the analysis of luck is the analysis of an ordinary concept that finds expression in ordinary language had better do everything possible to preserve our strongest intuitions regarding that concept in order to guard against substituting technical concepts different from the ordinary concepts we seek to analyze. Methodologically speaking, the severity of denying what Levy is denying—that the finding of buried treasure is lucky—should not be overlooked.⁷

⁶ As is the case when philosophers talk about luck, I assume Levy would understand talk of fortune to include both good and bad fortune, and would further add that, in ordinary talk, we often mean “good fortune” when we say “fortune”, and we use qualifiers like “bad” only when we want to make clear that we mean “bad fortune”.

⁷ Even Levy seems to worry that he might not have an intuitively well-motivated reason for rejecting that BURIED TREASURE is an instance of luck: “My intuition that Vincent is not lucky might be theory-driven, and might therefore simply beg the question against her” (2009, pp. 494). He goes on to offer some theoretical grounds which I consider shortly; for now, it is important to note that he thinks the fortune/luck distinction may not be easily drawn on intuitive grounds alone.

There is probably not a non-technical distinction to be drawn between fortune and luck. To see this, we should try to isolate the meanings of those terms. One way we can try to do this is to consider pairs of structurally identical sentences and try to ‘hear’ the differences in meaning between them. For instance,

“*The evening star* is visible today.”

vs. “*The morning star* is visible today.”

Contrasting this pair reveals that the two sentences express *something* different. One notices some difference in meaning between the two sentences. On the other hand, this pair reveals no difference:

“I was *lucky* to win that award.”

vs. “I was *fortunate* to win that award.”

It is easy to come up with more examples: just take any sentence where you are inclined to use “luck” or “fortune” and consider the pair of sentences which results from substituting the words. I doubt there are many, if any, ordinary sentences in English where “luck” could not be substituted with “fortune” without effecting a semantic change. That is a very compelling reason to resist the suggestion that speakers are confusing two similar concepts: what it suggests, on the contrary, is that we have two words for one concept.

Sometimes words mislead. But we now have strong *prima facie* evidence that luck and fortune are one and the same. I want to be clear about what I intend behind this test. I am not saying that this test is conclusive proof of semantic equivalence. What I am saying is that when two terms are inserted into structurally identical sentences and then read side-by-side, if no difference is heard, one thereby gains (at least) defeasible evidence of semantic equivalence. Further tests may reveal a difference, but some argument will be needed at that point to evaluate

whether the difference concerns pragmatic or semantic factors. For example, if I win an award and someone tells me that I am lucky to have done so, I may take it that the person has implicitly claimed that I do not deserve the award. My intuition on this point is not as strong if the person instead tells me that I was fortunate (rather than lucky) to win the award.⁸ But note that this further implication concerns what appears to be a pragmatic implication, not a semantic difference. (The sentences, “you were lucky to win that award” and “you were fortunate to win that award” sound, to me, like harsher and gentler ways, respectively, of saying the same thing.) I think that our test suggests strong evidence of semantic equivalence, and that defeating that evidence will require some plausible explanation of why we often confusedly attribute luck where we ought to attribute fortune instead.

Levy’s ‘confusion with fortune’ error theory attempts to provide just such an explanation.⁹ He claims that when we mistakenly attribute luck instead of fortune, it is because luck *is* very near in the vicinity: a fortunate event is a non-lucky event with luck in its proximate cause (2009; 2011).¹⁰ BURIED TREASURE is an instance only of *fortune*, Levy claims, because it is by *luck*

⁸ Thanks to an anonymous referee for encouraging me to clarify this point.

⁹ E.J. Coffman offers a different version of the ‘confusion with fortune’ error theory for all attributions of luck to events which occur in more than half of relevant nearby possible worlds. He claims that the winner of standard Russian Roulette isn’t lucky, and that we confuse the *appropriateness* of the assertion, “You are lucky you didn’t die in that silly ‘game’!” with its *truth* (2007). While it is appropriate to assert that sentence, Coffman argues, the assertion isn’t true. Someone making the assertion likely wishes to point out how irrational it is to play Russian Roulette in order to dissuade that person from doing it again. In order for this error theory to succeed, however, Coffman would need to convince us that when we attribute luck we are thinking about an imagined spoken sentence rather than a state of affairs. I’m confident that when I make or refrain from making a luck attribution in some instance, I’m judging that the concept LUCK applies in that case. I am not—or at least I need not be—imagining myself *saying*, “That was lucky.”

¹⁰ The appropriate nomological ‘distance’ between effect and proximate cause is left implicit in Levy. Because he is responding to Lackey, and because he does not challenge what Lackey might have in mind as the proximate cause(s), I take it that the proximate causes of Vincent’s finding of the treasure consists in something like the complex consisting of the following: Sophie and Vincent having the dispositions they have, the island having the geological and biological properties it has, Vincent’s mother having died, and Vincent’s knowledge of the nutritive requirements of rosebushes.

that Vincent and Sophie have the modally robust dispositions they have, and that state of affairs is a proximate cause of Vincent finding the treasure chest. The ‘kind’ of luck responsible here is *constitutive* luck: luck in one’s constitutive traits and dispositions (Levy 2009; 2011; following Nagel, 1979). Levy’s main reason for accepting that constitutive luck is a species of luck is that there is a strong intuition in favor of considering a child born with a congenital defect as unlucky: indeed, he claims that a view that denies that intuition “is misusing words, and leaves itself open to the objection that it avoids counterexamples by stipulation” (2009, pp. 495).

To make good on the error theory and provide a unified analysis of luck, Levy’s account needs to identify some genus of which chancy and non-chancy (e.g. constitutive) luck are species.¹¹ The account will also have to ensure that the two species-concepts are similar enough to each other that it is plausible that competent ordinary speakers would readily confuse them, misattributing luck to an event which isn’t lucky because another species of luck is causally responsible for that (non-lucky) event. A problem arises immediately: constitutively “lucky” traits are non-chancy, and so not obviously subsumable under the analysis of luck Levy defends. The very test Levy proposes for revealing whether or not an event is chancy—that of holding fixed some agents and temporally-prior events and varying other properties and events—reveals that the traits which constitute an agent vary only minimally in nearby possible worlds. The thought seems to be this: if S is a member of class F, and many Fs have trait C, then the nearby possible worlds

¹¹ Levy himself offers these success conditions for his analysis. He writes, “If there is an account of constitutive luck that is both plausible and not *ad hoc*, then constitutive luck is genuinely luck; the measure of *ad hocness* here is the degree to which the account diverges from [the account of chancy luck]” (2009, pp. 496). If my argument in the next few paragraphs is successful, then the two accounts of luck differ not in degree, but in kind. Thus the account of constitutive luck is *ad hoc*.

which include S will usually include S's being C. But if that so, then S's being C is not chancy, but *probable*.¹²

How, then, would ordinary speakers confuse Vincent's being lucky with his merely being fortunate if his being fortunate is a result of him having some trait he would probably have had in similar circumstances? Since Levy is defending an error theory that says speakers confuse luck with fortune, Levy has to find some luck in the proximate causes of Vincent's finding the treasure. For this, he proposes that Vincent's behavior was grounded in a trait for which he is constitutively lucky, rather than chancily lucky. Of course, that is not enough to make plausible Levy's claim that speakers readily *confuse* (chancy) lucky events with events that are non-lucky but have (constitutive) luck in their proximate causes. The problem is that if the concepts are not similar enough, the claim that speakers confuse the two will not be plausible.¹³ Levy attempts to subsume chancy luck and constitutive luck under the genus 'luck' by defining non-chancy luck so that it includes the significance and lack of control conditions, but has a different improbability condition. Levy claims "there are some close parallels between constitutive luck and chanciness" (2009, pp. 496). Having given up analyzing constitutive luck using a modal improbability condition, he instead opts for a statistical improbability condition: traits that vary widely within some (relevant) reference class¹⁴ of which I am a member are traits for which I am constitutively lucky (2009, pp. 496; 2011, pp. 33).

¹² See footnote 11 for discussion of the problems raised for Levy by this understanding of luck.

¹³ This follows from the problem of how constitutive and chancy luck are related plus Levy's claim that speakers confuse lucky events with fortunate events. My argument is that if the analyses of fortune and luck Levy proposes were the ordinary concepts, the confusion would be rather surprising because speakers would be confusing two very different things.

¹⁴ The view faces the ever-present problem of the reference class which faces statistical analyses generally. Levy is aware of the problem, and claims that the relevant reference class is fixed by context (2011, pp. 34). It is less clear whether he thinks the appeal to context actually solves the problem: a person who survives a plane crash is lucky *qua survivor of a plane crash* to have only lost three limbs; the very same person *qua human being* is very unlucky

The analysis of luck Levy is offering calls upon an improbability condition, and that condition is spelled out using two distinct concepts of probability. Introducing these into the analysis vitiates the attempt to provide an analysis of a single, unified concept. The new analysis of luck is disjunctive: (given the other conditions) an event is lucky if EITHER it is the result of a trait which is statistically unlikely within a relevant reference class OR the event fails to occur in a large enough proportion of nearby possible worlds with the same relevant initial conditions. Put a little differently, if Levy is trying to provide an analysis of a concept, the improbability condition must state, “E is either statistically improbable or modally improbable.” That the analysis includes as a single necessary condition the satisfaction of one or the other of two *distinct* concepts of probability entails the two concepts are distinct.¹⁵

The disjunctive analysis does not serve Levy’s purposes well. First, one of the motivations behind offering the ‘confusion with fortune’ error theory was to explain why we might confuse fortune with a *closely related* concept, namely, the concept of luck. But once the analyses are spelled out, these do not appear to be closely related: it is unclear how fortune and luck are appropriately similar because it is unclear how constitutive luck and chancy luck are appropriately similar. The two concepts of luck share, in rough description, a few necessary conditions—significance and lack of control—but the concepts rely on notions of probability very distinct. It seems unlikely that even ordinary speakers are confusing modal probability with statistical

to have done so. In ordinary discourse context by itself does not solve the problem of the reference class: we often use, in addition, some clarifying device to indicate the reference class we have in mind, e.g. “So many of the people on the flight with you died; you’re lucky to have only lost some limbs.”

¹⁵ Suppose Levy were to drop the idea that we misattribute luck to probable events because there is luck in the non-lucky event’s proximal cause, and keep part of the error theory otherwise intact, claiming that we misattribute luck when a statistically improbable event or trait is a proximate cause of some event. This trimmed-down error theory would avoid the complications related to the analysis of luck discussed above, but it would also lose all plausibility as an error theory: the mistake speakers make when luck is misattributed would then be locating luck where there isn’t any to be found at all.

probability. Ordinary speakers have the resources to distinguish between the two concepts: we talk about things that could have easily been different (modal probability) and characteristics that aren't common (statistical probability), and we understand the difference between these notions. Even if one thinks ordinary speakers lack the discernment to spot the difference, certainly philosophers—including those who find BURIED TREASURE compelling—do understand the difference. Second, another main motivation behind offering the error theory was that there seemed to be too much luck. But now, there will be too much fortune. Any event which is the result of the activity of a statistically unlikely trait in some agent is fortunate. My own linguistic intuitions are quite content with positing an abundance of luck; why should Levy's be better served by positing an abundance of fortune?

Levy was compelled by the intuition that an event could be modally probable while still being lucky. In order to capture that intuition but avoid threatening the defensibility of the modal account, he offered a 'confusion with fortune' error theory. I have argued that it is not possible to both defend a modal improbability condition and respect our intuitions about probable lucky events. I propose to simply move forward with a non-modal improbability condition.

The Epistemic Analysis of Luck

Thus far I have argued that modal improbability is unnecessary for luck. I want to argue that a different notion of probability is at work in luck. To start, consider the way in which our attributions of significance to subjects behave. For an event to be lucky requires that it significant to the subject for whom an event is lucky. As I noted briefly above, just what analysis of

significance ought to be given is debatable.¹⁶ But any plausible account of significance will allow that whether an event is significant can vary by subject. If E matters to you but not to me, then E could be lucky for you, but not for me. Take a subject for whom an event is improbable enough to be lucky and for whom the event is very significant. That event is lucky for that person. Now substitute someone who is identical in every way except that for the substitute, the event is insignificant.¹⁷ The event is non-lucky for that person. This result follows from the significance condition on any account of luck, regardless of the kind of probability required by the account.

On the view I defend, this same sort of variation applies to the satisfaction of the improbability condition, not just the significance condition. The improbability condition I defend is an *epistemic* improbability condition. One and the same event can have different probability assignments for different epistemic subjects, and therefore one and the same event can have different luck assignments (even when the significance of the event is equal to the two subjects).

The intuitions we surveyed early in the paper lead naturally to an epistemic account of luck. Lucky events are often hard to predict and surprising, as are events for which I lack sufficient evidence to anticipate. An epistemic improbability condition easily captures the intuitions that seemed to support the modal accounts. Our intuitive judgments regarding what happens in nearby possible worlds are guided by our appreciation of the relevant evidence. We judge that an event

¹⁶ My own view, which I have no space to defend in this paper (but which I hope to defend at a later time) is that an event is significant to a person when that person would regard the event as significant, when informed about the event and the conditions surrounding the event. Notice, first, that this is not a definition of “significant”, because “significant” appears in the definiens. Second, this counterfactual subjective significance condition allows for an event to be lucky for a subject when (1) the improbability condition is met and (2) the subject does not *actually* know about the circumstances leading up to the event: even, perhaps, that the event has already taken place.

¹⁷ Someone who thinks that significance should be understood as an objective notion cannot allow that in a case like this, a duplicate subject with different preferences from the original subject can have varying significance-attributions. That is because (presumably) what is objectively important for any subject is precisely what is objectively important for any individual who is identical to that person, regardless of the person’s attitudes.

would or would not happen in some proportion of nearby possible worlds on the basis of experience. Experience gives us a sense of what *would* happen in certain circumstances in virtue of what *has* happened in similar circumstances: an event which, *on our evidence*, is likely to happen is one that we (sometimes) judge happens in a large number of nearby possible worlds. The intuitive judgments driving the modal account have their origin in assessments of evidence.¹⁸ For that reason, the epistemic account can successfully mimic the intuitive verdicts on cases that the modal accounts deliver. So, an epistemic improbability condition can deliver the verdicts on cases we seek, and it can mimic any verdict that the modal theorist can offer by appeal to possible worlds.

To my knowledge, the only other person to argue for an epistemic interpretation of the improbability condition on luck is Asbjørn Steglich-Petersen (2010). He defends the following improbability condition: “*S* is lucky with respect to *E* at *t* only if, just before *t*, *S* was not in a position to know that *E* would occur at *t*.” As he recognizes, what it takes to be “in a position to know” varies by philosopher. He claims that one’s knowing that an event will occur entails having all other positive epistemic statuses with respect to the proposition which says that the event will occur (2010, pp. 369). Depending on one’s epistemological views, however, this claim is false. Take G.E. Moore, for instance. He thought that it was possible to know that pencils exist. He did not think that his knowing this entailed that he had evidence which guaranteed for him that there

¹⁸ Note that the defender of a modal account cannot reciprocate my claim about the priority of the epistemic to the modal. Although it seems that there is an intuitive sense in which there could be a higher or lower ‘proportion’ of near and relevant possible worlds in which some event happens given some (to-be-specified) initial conditions, there are, presumably, *infinitely* many near and relevant possible worlds in which the event occurs. So the modal theorist requires some way of relating sets of infinitely many possible worlds in order to make sense of these ‘proportions’. I don’t know what the solution might be for the modal view. It looks like a serious problem to me. (Cf. Fumerton 2014).

were pencils. So, if Steglich-Petersen were to ask Moore if it was at all lucky, in any relevant sense, that his belief about pencils is true, Moore would say “no”, even while it is epistemically possible for Moore that his belief is false.¹⁹ But I think that Steglich-Petersen wants Moore’s belief to turn out at least a little bit lucky. It is not clear, however, how the account will deliver that result. Even if Steglich-Petersen’s account could be stated in such a way that it gave that verdict, it seems to me that the account will do best to discuss epistemic probability, and not knowledge. For that reason, I will avoid reference to knowledge in my analysis of the improbability condition on luck.²⁰

Luck comes in degrees. I defend the following definition of the improbability condition on luck:

(EAL) Event E is lucky to some degree for subject S when E was/is not epistemically guaranteed for S.

I call this the Epistemic Analysis of Luck (EAL).²¹ I take the epistemic guarantee of an event to be equivalent to the event’s having an epistemic probability of 1 on one’s evidence. Furthermore, how lucky an event is is partly a function of how epistemically probable the event is. Holding fixed the significance of some event, if the event has an epistemic probability of 1, then the event is not at all lucky, to any degree whatsoever; if the event has an epistemic probability of .95, then it was a little bit lucky; if the event has an epistemic probability of .1, then it was very lucky; and so on.

¹⁹ Moore is clearest about this in (1959). He there seems to hold that he can be certain that p is true even while he does not have evidence guaranteeing the truth of p.

²⁰ All of that being said, however, I believe that Steglich-Petersen and I are much in agreement about the analysis of luck, at least with respect to the improbability condition.

²¹ In my view, but not explicitly in this paper, EAL refers to the conjunction of this condition plus a counterfactual subjectivist significance condition. That condition says E is significant for S if the subject would, upon some amount of reflection, care whether E. Defending that condition on EAL will require a separate paper.

This account seems to suggest that there is quite a bit of luck in the world (so long as the significance condition is also met). If the condition claimed only that a lucky event is one that is epistemically *improbable*—that is, has an epistemic probability of 0.5 or lower—and significant, then the account would be open to the same sorts of counterexamples as the modal accounts. It is epistemically probable that I will win a game of twenty-chamber Russian Roulette.²² But winners of that game are lucky, so the epistemic improbability condition is insufficient. There are only two options: specify an arbitrary threshold below which events are epistemically improbable enough to satisfy the improbability condition, or allow that any event is lucky to some degree when its occurrence is not epistemically certain for the lucky subject. At risk of repetition, there is intuitively a lot of luck in the world unless much of our luck-talk is false. That is why the condition is stated as requiring only epistemic *possibility*: the lack of an epistemic guarantee. These probabilities are determined by epistemic factors, not psychological ones. Psychological certainty is not epistemic certainty. Simply being absolutely (but unjustifiably) *sure* that an event will occur does not make the event non-lucky.

Let us return to the cases that got our discussion started. EAL delivers the correct verdicts on BURIED TREASURE and RUSSIAN ROULETTE. In both cases, the events are significant to the subject. In the former, EAL claims that Vincent is lucky because the epistemic probability for him that he would find a treasure chest was very, very low. In the latter, Samuel is lucky to win because he couldn't be epistemically sure that the chamber with the bullet wasn't next in line to fire.

²² Plausibly, the *a priori* probability of a Russian Roulette win is 5/6. That is why a strict epistemic *improbability* condition does not, by itself, seem to fare better than a modal condition.

EAL is open to various accounts of how epistemic probabilities are determined. My own view is that epistemic probabilities are determined by features which are reflectively accessible to a subject, namely a subject's experiential states and beliefs, plus the evidential relations between these states and beliefs. But that is not the only account of how epistemic probabilities are determined, nor is it the only account consistent with EAL. For example, one might hold that it is possible that one's evidence can consist at least partly in facts that guarantee the truth of one's beliefs, even for propositions about the external world. Then one could hold EAL and claim that we have lots of beliefs about the external world that are true and not at all lucky, because we possess evidence which guarantees that they are true. Therefore, EAL alone does not guarantee that intuitively non-lucky events are in fact lucky: only EAL combined with an account of epistemic probability can do that. That being said, I do not wish to disguise the fact that the way I think of EAL, many events that one might not readily *say* are lucky are, in fact, lucky, so in the following discussion I will assume that we are combining EAL with an account of evidence which allows that we often lack such factive evidence.

It will help EAL be viewed in a positive light if I can argue that we can divorce the truth of luck-attributions from their appropriateness. I said above that, pragmatically speaking, luck attributions seem to undermine the attribution of credit and responsibility: if I say you were lucky to win some award, you might take offense that I seem to have implied that you did not deserve to win it. Now, I do not think that we naturally *say* that highly (modally or statistically) probable events are lucky. For example, according to EAL, Tiger Woods is at least a tiny bit lucky whenever he hits a ball beyond 50 yards. Since it was epistemically *possible* for Woods that the ball would fail to go 50 yards, then on the assumption that it is significant to him that the ball go further than 50 yards, Woods is the slightest bit lucky that his drive goes as planned. It would definitely not

be appropriate to *say* so, because that would pragmatically imply that Woods does not deserve credit for his successful drive. Still, the hit is just a *little* lucky. His club head could have broken off, a pigeon could have flown into the ball at the 10 yard mark, he could have had a heart attack during his backswing, etc. Once possibilities like these are raised, it is easy to admit that Woods' long drive is a little lucky. However, it would be wildly inappropriate for an onlooker to see Woods' drive and assert, "That was a lucky shot!" The improbability of his drive passing the 50 yard mark is so miniscule as to be unassertable.

Even when we would not be too naturally inclined to attribute luck to a likely event, making various possibilities salient causes us to recognize that the event was not as sure to happen as we might previously have claimed, and thereby to attribute some degree of luck. When we raise these sorts of possibilities, even about events that we continue to believe were overwhelmingly likely to occur, we recognize that some amount of luck was present. In this way, our practice of luck attribution behaves much like our practice of knowledge attribution: it sounds as bizarre to say, "I *know* that P, but I'm *not sure* whether P," as "It's not *at all* lucky, *to any* degree, that E, but I *didn't* have *conclusive* reasons to expect E".²³ This remarkable similarity suggests that there is something in common between luck- and knowledge-attributions: if EAL is true, and luck-attributions are assessments of epistemic probability, that would explain the similarity. But notice that the comparison only follows for an account of knowledge which connects knowing a proposition with that proposition having an epistemic probability of 1.²⁴

²³ Just how to understand knowledge attributions and retractions like these is a matter of present debate. Defending EAL does not require taking a position. For representative positions, see DeRose 2009, Rysiew 2001, Stanley 2005, Fantl and McGrath 2009, Unger 1975.

²⁴ See Williamson 2002 for an attempt to accept this without accepting skepticism.

When we attribute luck to a subject, we often do so by taking up that subject's perspective. From the perspective of an omniscient being nothing is (or would be) lucky, on the assumption that omniscience requires knowledge of future contingent propositions, because the truth value of every proposition would have an epistemic probability of 1. It doesn't follow from RUSSIAN ROULETTE being non-lucky from God's perspective that it is in fact non-lucky. Earlier, I argued that any plausible account of the significance condition will allow that a single event could be significant for one person, and thus a candidate for luck for that person, while insignificant for another, and consequently not a candidate for luck for that person. Evaluating the proposition *Samuel is lucky to be alive* depends upon selecting an epistemic perspective from which to evaluate the proposition. When I say that Samuel is lucky to be alive, I could mean that Samuel's surviving was not epistemically certain *for me*. Or I could mean that Samuel's surviving was not epistemically certain *for him*. Evaluating the truth of a luck-attribution requires specifying the epistemic position from which we are making the attribution. So, according to EAL, whether an event is lucky requires (1) evaluating the event for significance from different perspectives, and (2) evaluating the event for epistemic probability from different perspectives.

This completes my presentation of EAL. I now consider two important objections.

Odd Luck

EAL claims that an ordinary event which is epistemically possible for some person and which is significant for that person is lucky. In the case of low epistemic probabilities, the events are not only lucky, but highly lucky. Suppose there are persons who, for some reason or other, have epistemic probability assignments out of step with most of ours. Then it may follow that because of their bizarre epistemic probabilities, the truth of luck-attributions from their perspective varies

from the truth of luck-attributions from other perspectives, including ours. This is the odd luck objection.

This objection is not worrisome for EAL. First, epistemic probabilities are constrained by rules on every plausible account of epistemic probability. That a probability is epistemic requires, in part, that it be conducive toward truth. The “oddness” of these individuals’ probability assignments therefore cannot be due to simple wishful thinking or cognitive vice. Epistemic probabilities for a subject are not open to tampering independently of evidence possessed by that subject. So, for a problematic odd luck scenario to obtain would require that a person have a bizarre set of experiences which actually epistemically justify propositions to a degree different from the rest of us. Any plausible theory of epistemic probability will hold that it is possible for a proposition to be justified to some degree for some individual while that same proposition has a different probability for some other person. The existence of unusual believers like these does not impugn EAL.

Second, any subjectivist significance condition entails a similar result. Some individuals take more events to be significant, and to a higher degree. When this happens, subjectivist accounts of significance allow that these persons are better positioned with respect to how much luck they will enjoy.

Luck and Ignorance

Because EAL makes the possession of luck dependent upon epistemic probabilities, what happens when a subject lacks any reason to suspect that some intuitively lucky event has/did/will occur?

The following case illustrates this problem.

KILLER COKE: Unbeknownst to me, all cans of Cherry Coke in my region have been accidentally poisoned as the result of a shipping factory mishap. The only

person aware of the mishap is the Regional Director, who learned of the mishap too late to stop the cans from being sold, but before anyone drank any of the hazardous beverages. I walk to my refrigerator, open a can, and take a sip. I don't die: also unbeknownst to me, another mishap occurred when a can of Cherry Coke owned (but never opened) by a driver from a different region accidentally made its way into my region's shipping facility and ended up in my twelve-pack. That's the one I'm drinking now. Thanks to this series of accidents, I got lucky.²⁵

Intuitively, one might think, my not dying from sipping that Cherry Coke is lucky. But the epistemic probability for me that *this* can would be a killer was no higher than it usually is: that is, it is low.

EAL is not defeated by this objection. First, when we consider the proposition *surviving the killer Coke was lucky* we are considering an ambiguous proposition. Luck-attributions are made and evaluated from an epistemic perspective, and we are not always explicit about which perspective we are adopting. In this case, as in others, we have to choose the perspective from which we will evaluate the luck-attribution. The perspectives that are relevant for this case are the Regional Director's actual epistemic perspective, the Regional Director's counterfactual epistemic perspective, or my perspective. The Regional Director could truly and warrantably assert "that was lucky", so long as the proposition he has in mind is (i) *I (the Regional Director) am lucky that E* or (ii) *if I (the Regional Director) were the person drinking the Coke, I would be lucky*. His assertion is clearly true on interpretation (i): at a minimum, he would have lost his job for unwittingly releasing poisonous beverages for general consumption. And on interpretation (ii), he would be considering my brush with death from his more informed perspective and would therefore regard the event as lucky if he were me. Only if what the Regional Director wishes to assert is the proposition *the Coke-Drinker is lucky from the Coke-Drinker's own perspective* does

²⁵ Thanks to Richard Fumerton for this objection.

his assertion even threaten being problematic. Even in that case, though, the Coke Drinker *is* lucky from the Coke Drinker's perspective, because the significance condition is met (were the Coke Drinker in possession of evidence concerning the events leading up to drinking the Coke, the Coke Drinker would undoubtedly regard the event as significant) and so is the improbability condition (the proposition *this Coke is non-lethal* was not epistemically guaranteed, and never is). The only potential worry now is that the proposition *this Coke is non-lethal* is no more epistemically improbable for the Coke Drinker now than it normally is.²⁶ So it may appear that EAL delivers the wrong result on *how* lucky an event is. The appearance is illusory: the probabilities being what they are, the event is far more significant to the subject than the drinking of a normal Coke is, because the subject would regard this situation as more significant than other situations in which the subject drinks a Coke. How significant an event is partly determines how lucky it is. That being the case, it is at least unclear that KILLER COKE poses a serious objection to EAL.

Second, those still worried by this case may wish to replace the epistemic improbability condition stated above with a counterfactual epistemic possibility condition, so that the degree to which an event is lucky in the actual world varies partly with how epistemically probable the event *would* be for that subject under certain counterfactual conditions. One way to articulate this modified version of EAL would be to tie the counterfactual epistemic probability condition to the counterfactual significance condition: the account would then state that for purposes of luck-attribution, the probability of the event is equal to how probable the event is in whatever conditions need to be realized for the subject to count the event as significant. Further qualifications are clearly needed, however: once the subject has the evidence necessary to make an informed

²⁶ Thanks to an anonymous referee for helping me to clarify this point.

significance judgment, the epistemic probability of the event has also increased considerably, and therefore will count as very significant, but also very probable, and consequently not lucky. Were I to consider the sequence of events that led to my drinking a safe Coke, I would indeed find it significant. (I now know that I almost *died*, after all.) For the same reason, at the time of your asking me, I acquired testimonial evidence that I narrowly escaped death. Maybe the counterfactual epistemic probability condition should therefore take the epistemic perspective of an informed actual-world third-party, plus the already stated counterfactual subjective significance condition. The difficulty now is choosing among unequally informed (and possibly non-actual) third parties. Perhaps an account of a counterfactual epistemic possibility condition could be developed which would avoid the unwelcome results just described. I find no reason to develop one here, however, because I find the first reply satisfactory: once we disambiguate the possible meanings behind the assertion, it is clear that the luck-attribution is true from every perspective including the Coke drinker's, and therefore we are able to make perfect sense of our luck attributions being not only appropriate, but also true.

References

- Coffman, E.J. "Thinking about Luck", *Synthese*, 158 (2007) 385-398.
- DeRose, Keith. The Case for Contextualism. Oxford University Press, 2009.
- Fantl, Jeremy and Matthew McGrath. Knowledge in an Uncertain World. Oxford University Press, 2009.
- Fumerton, Richard. Knowledge, Thought, and the Case for Dualism. Cambridge University Press, 2014.
- Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis*, 23 (1963) 121–123.
- Grice, H.P. "Logic and Conversation", Syntax and Semantics, vol. 3, eds. Peter Cole and Jerry L. Morgan, pp. 41-58. Reprinted in A.P. Martinich, "The Philosophy of Language", fifth edition, Oxford University Press, 2008.
- Hanna, Nathan. "Moral Luck Defended", *Noûs*, forthcoming.
- Lackey, Jennifer. "What Luck Is Not", *Australasian Journal of Philosophy*, 86:2 (2008) 255-267.
- Latus, Andrew. "Constitutive Luck", *Metaphilosophy*, 34:4 (2003) 460-475.
- Levy, Neil. Hard Luck. Oxford University Press, 2011.
- . "What, and Where, Luck Is: A Response to Jennifer Lackey", *Australasian Journal of Philosophy*, 87:3 (2009) 489-497.
- Moore, G.E. "Four Forms of Skepticism" in Philosophical Papers. George, Allen and Unwin, London: 1959.
- Nagel, Thomas. "Moral Luck", Mortal Questions. New York: Cambridge University Press (1979) 24-38.
- Pritchard, Duncan. Epistemic Luck. Oxford University Press, 2005.

- Rescher, Nicholas. Luck: The Brilliant Randomness of Everyday Life. Farrar Straus Giroux, 1995.
- Rysiew, Patrick. "The Context-Sensitivity of Knowledge Attributions", *Noûs*, 35:4 (2001) 477–514.
- Stanley, Jason. Knowledge and Practical Interests. Oxford University Press, 2005.
- Steglich-Petersen, Asbjorn. "Luck as an Epistemic Notion", *Synthese* 176 (2010) 361-377.
- Stevenson, C.L. "The Emotive Meaning of Ethical Terms", *Mind* 46 (1937) 14-31.
- Unger, Peter. Ignorance: A Case for Skepticism. Oxford University Press, 1975.
- Williamson, Timothy. Knowledge and Its Limits. Oxford University Press, 2002.