LUCK, KNOWLEDGE, AND EPISTEMIC PROBABILITY

Gregory STOUTENBURG

ABSTRACT: Epistemic probability theories of luck come in two versions. They are easiest to distinguish by the epistemic property they claim eliminates luck. One view says that the property is knowledge. The other view says that the property is being guaranteed by a subject's evidence. Asbjørn Steglich-Petersen defends the Knowledge Account (KA). He has recently argued that his view is preferable to my Epistemic Analysis of Luck (EAL), which defines luck in terms of evidential probability. In this paper, I defend EAL against Steglich-Petersen's arguments, clarify the view, and argue for the explanatory significance of EAL with respect to some core epistemological issues. My overall goal is to show that an epistemic probability account of luck rooted in the concepts of evidence and evidential support remains a viable and fruitful overall account of luck.

KEYWORDS: luck, epistemic probability account of luck, knowledge

1. Two Epistemic Probability Accounts of Luck

> You stop for lunch at your favorite hot dog stand near your office. The owner says, "You come here all the time! Here's one for free."

By luck, you got a free hot dog. It is natural, as a first pass, to think of yourself as lucky to receive your free lunch because you had no idea that you would. You are surprised, you had no reason to expect this, you did not know it would happen. These general thoughts support an *epistemic* theory of luck: lucky events are those significant events that you lack good epistemic reasons to expect.

Here are two different ways of being slightly more precise about this case. One way: you did not *know* that by going to the stand today, you would get a free hot dog (and per hypothesis receiving the hot dog is significant to you). Another way: the *evidence* in your possession did not guarantee that you would receive a hot dog (and per hypothesis receiving the hot dog is significant to you). What both explanations have in common is that *epistemic* factors are used to explain why this event is a lucky one. What sets them apart?

Distinguishing those two superficially similar explanations requires taking positions on several further issues concerning the nature of knowledge, the nature of evidence, what it is to possess evidence, and how evidence can support a proposition sufficiently to 'guarantee' it. Both, though, are probability theories of luck that identify *epistemic* probability as the kind of probability that matters for luck. There is no need to rehearse different categories of views of luck and pointing to the important similarities between the knowledge account and the evidential probability account, or in rehearsing how these two epistemic accounts of luck compare with non-epistemic accounts.[1] This paper is about how exactly to characterize the epistemic properties that factor into luck. This is a place for highlighting *differences*.

Steglich-Petersen defends this view, which he calls the "knowledge account:"[2]

> KA: 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.

I have defended an alternative view, and called it the "epistemic analysis of luck:"[3]

> EAL: S is lucky with respect to E only if the evidence S had immediately prior to E occurring did not guarantee that E would occur.[4]

The fundamental concept in KA is the concept of *being in a position to know* that an event will occur. In EAL, it is the concept of *a subject's evidence guaranteeing* that an event will occur. The difference between those two concepts is more important than one may think, even while the views have much in common. Steglich-Petersen has argued that KA is superior to EAL, as only KA can handle a certain class of cases. The chief failing of EAL, according to Steglich-Petersen, is that there are conditions on knowledge that make a difference for luck, and these conditions are neglected by EAL, with its focus on evidential probability as the sole epistemic condition relevant to luck.

This paper proceeds as follows. Section 2 evaluates an argument that EAL cannot account for lucky necessities and solves the problem by appealing to intuitive conditions for evidence possession that are not met in cases of lucky

---

[1] Cf. Fernando Broncano-Berrocal, "Luck," *The Internet Encyclopedia of Philosophy*, https://www.iep.utm.edu/, 2020; Steven D. Hales, "Why Every Theory of Luck is Wrong," *Nous*50, no. 3 (2016): 490-508; Gregory Stoutenburg, "In Defense of an Epistemic Probability Account of Luck," *Synthese* 196, 12 (2019): 5099-5113.

[2] Asbjørn Steglich-Petersen, "Luck as an Epistemic Notion," *Synthese* 176, 3 (2010): 361-377, Asbjørn Steglich-Petersen, "Does Luck Exclude Knowledge or Certainty?" *Synthese* (2018), https://doi.org/10.1007/s11229-018-1790-z.

[3] Gregory Stoutenburg, "The Epistemic Analysis of Luck," *Episteme* 12, 3 (2015): 319-334, Stoutenburg, "In Defense."

[4] Steglich-Petersen calls this the 'certainty' account, but that label unfortunately leaves aside the concepts of evidence and evidential probability that are central to EAL. EAL also requires that the event is significant to the subject.

necessities. Section 3 argues that a higher-order evidence version of the lucky necessities case does not undermine EAL because higher-order evidence at least sometimes has lower-order import. Section 4 argues that even apparently non-probabilistic evidence can be understood probabilistically. Section 5 clarifies the role of evidence in EAL. Section 6 argues that EAL, and not KA, has the resources to provide an interesting, unified account of epistemic luck that explains many common intuitions about knowledge, including intuitions about closure principles, Gettier cases, and lotteries.

2. Logical Necessities, Knowledge, and Evidence

Steven Hales maintains that every theory of luck is open to counterexample.[5] One of his counterexamples is this:

> Logical Bandit: The logical bandit points a gun at you and tells you that unless you correctly answer a logic puzzle, he's going to steal your wallet. He gives you this poser:
>
> Suppose you go to a diner where the cook is famous for pancakes. Actually, he is famous for burning 50% of the pancake-sides he cooks, and cooking the other 50% perfectly. The statistics: One third of his pancakes are golden on both sides; one third are black on both sides; and the remaining third are golden on one side and black on the other. You order a pancake. When it comes, the side you can see is golden. What is the chance that the other side is golden?
>
> You are horrible at this sort of thing, and are completely flummoxed by the gun, the puzzle, and whole situation. You make a wild guess and say "it's 2/3." The logical bandit, who could tell you were just guessing, smiles ruefully and replies, "you're lucky the correct answer is indeed 2/3," and vanishes into the night.[6]

The point of the counterexample is that the answer 2/3 is necessary, but all modal and probability theories of luck say that an event is lucky only if it could have failed to occur (or could have failed to be true) or has a probability below 1.0.[7] So, concludes Hales, all theories of luck, including probability theories, are false.

I have argued that EAL can handle the Logical Bandit:

> You do not understand the puzzle, so the details of the puzzle that entail the correct answer are not included in your evidence. If the probability of the answer being 2/3 had been 1.0 on your evidence, you would have been in a position to know that 2/3 was the answer, and consequently not lucky to answer correctly.

---

[5] Hales, "Why."

[6] Hales, "Why," 495.

[7] Intuitively, an event counts as *im*probable only if its probability is below 0.5. The requisite notion of improbability for luck is technical.

> You clearly were not in a position to know the answer, since you had to guess. So, the probability condition is satisfied.[8]

Since the subject does not know that the answer is 2/3, that implies that the evidential probability for the subject is below 1.0. Steglich-Petersen objects: "On the face of it, [the probability that 2/3 is the answer] is 1, since it is entailed by the details of the puzzle."[9]

It is important to separate the question of how probable it is that the answer to the puzzle is 2/3 from the question of how probable it is *on the subject's evidence* that the answer to the puzzle is 2/3. The questions are different. EAL says that the probability that determines luck is *evidential* probability, so the probability of the answer to the puzzle being 2/3 does not matter for luck. What matters is whether the probability that the answer is 2/3 is 1.0 *on the subject's evidence*. If and only if it is, then Steglich-Petersen is right that the evidential probability account fails against the case.

Steglich-Petersen says, "On the face of it, [the probability that 2/3 is the answer] is 1, since it is entailed by the details of the puzzle."[10] The reasoning seems to be this: The Logical Bandit tells S the details of the puzzle, so S has the details of the puzzle as evidence, and those details yield a probability of 1.0 that the answer is 2/3. For that reasoning to be sound requires the truth of this principle, or at least something very close to it:

> Hearing Implies Having (Hearing): If S hears and understands that p, then p is included in S's total evidence.

Without (Hearing) or something much like it, there is no clear reason why we should think that the Bandit stating the details of the puzzle puts the details of the puzzle into the subject's evidence.[11]

There are examples that support (Hearing). I ask the gate agent if it will be a full flight. The agent says "no." It is reasonable to include *it is not a full flight* in my total evidence, in the set of propositions and experiences that I rely on. Ordinary cases of testimony are like this. But (Hearing) is not true. If the agent answers "no"

---

[8] Stoutenburg, "In Defense," 10.

[9] "Knowledge or Certainty," 6.

[10] "Knowledge or Certainty," 6.

[11] Similar principles could be invented for other sensory modalities, and similar concerns would arise. Does a brief memorial flash of an image in my mind put into my evidence facts about a distant family event, if the event was the original cause of the image? Does the faint smell of burning metal put into my evidence that I am in a motorcycle factory, if I am in a motorcycle factory and should be able to figure that out? The suggestion in each parallel case is that the threshold for what makes something in experience count as evidence is beyond mere exposure.

and I am immediately distracted by something else, to the point that the interaction is forever beyond retrieval for me, then the proposition *it is not a full flight* will *not* be included in my evidence. Hearing and understanding that p is insufficient for having p as evidence.

The relevant notion of "understanding" is critical here, too. Even without a full account of understanding, we need a relatively clear idea of what role understanding plays in turning an auditory sensation into evidence. A person who can repeat spoken sentences from another language does not count as understanding the language unless that person can also perform certain inferences like correctly processing alterations in word suffixes by number or gender and is able to answer at least some questions about what was spoken. Someone who can repeat the sentences but cannot do these other things is merely parroting words, without understanding.

The question now is, is the Logical Bandit case more like a straightforward instance of receiving testimony, or is it more like an instance of hearing something and immediately forgetting it, or hearing sentences from a foreign language? If the details of the puzzle are included in the subject's evidence, and they are all the relevant evidence the subject has concerning the answer to the puzzle, then EAL would indeed imply that the probability for the subject that 2/3 is the answer is 1.0.

I claimed that the hearer in the Bandit case does not understand the details of the puzzle, which implies that the subject does not have the details of the puzzle as evidence.[12] Steglich-Petersen objects that "there is no reason why he should not [understand the details of the puzzle]" and in support of the claim points out that the details of the puzzle are simple, before claiming that the subject merely struggles to draw the correct inference from the puzzle. Here, though, the subject makes a *complete* guess. It is not just that the subject fails to draw the correct inference, but that the subject fails to draw any inference *at all* and instead must resort to picking a number at random. That is a good reason to think that the subject does not understand the puzzle because the subject is in a position like that of a person who is hearing sounds from another language. There is of course a distinction between failing to infer and failing to understand, but in this case being able to competently think about how one might devise an answer to the question is constitutive of understanding the puzzle, just as having competence with a language is needed if hearing spoken words in that language puts the words spoken into one's evidence.

[12] "In Defense," 10.

So, it is reasonable to say that the subject indeed does not understand the puzzle. And thus, the details of the puzzle are not included in the subject's body of evidence from which the probability of the answer being 2/3 can be derived.

3. Evidence and Higher-Order Defeat

Steglich-Petersen offers a variation on the Logical Bandit intended to show that if the subject *does* understand the puzzle, EAL gets the case wrong for another reason. The idea is that evidential certainty is incapable of eliminating luck, because it is possible for p to be certain on one's evidence while p is nevertheless lucky. The new case is a modification of the Logical Bandit, with a twist concerning higher-order defeat. The case starts the same as before, then continues:

> Logical Bandit with Higher-Order Defeat: You are usually pretty good at logic puzzles, and after a bit of thinking, you reach the correct answer of 2/3. However, the bandit then tells you that he slipped a powerful reason-distorting drug into your coffee just before pointing his gun at you. Reasonably, this convinces you that you have most likely made a mistake, even though in this instance you did not in fact make a mistake.[13]

Yet the subject guesses 2/3 and is intuitively lucky that the answer is so. Steglich-Petersen says that the luckiness here cannot be explained as a matter of epistemic probability, because the first-order evidence that 2/3 is the answer is the same now as it was in the first version of the case.

In arguing for his verdict, Steglich-Petersen says of the subject's guess: "*In his evidential situation*, this answer isn't much better than any other…"[14] But that explanation plainly *supports* the EAL explanation of the case. On the subject's *total evidence* the probability that the answer is 2/3 is indeed *under* 1.0, because in the new version of the case the subject's evidence includes the details of the puzzle, which taken alone would make the probability of 2/3 equal to 1.0, *and* the additional evidence that the subject's reasoning is faulty, which *lowers* the probability that the answer is 2/3. Higher-order evidence is evidence, so it affects epistemic probability.

An intuitive way of reaching the same conclusion is this. Take the subject immediately after the subject has arrived at the answer of 2/3, and ask: "How likely is it that the answer is 2/3?" Presumably, the subject will be very highly confident that the answer 2/3 is true, and will perhaps say: "It is definitely 2/3." Now, let the Bandit make the claim about the coffee, and ask again: "Now that you have heard

[13] "Knowledge or Certainty," 7-8.
[14] "Knowledge or Certainty," 8.

that the Bandit drugged you, how likely is it that the answer is 2/3?" Here, presumably the subject will give some answer that is lower than the initial answer. If the higher-order evidence is irrelevant to epistemic probability, the probability of the 2/3 answer has not gotten *even a little* lower than it was prior to being told that you have been drugged. It has stayed *exactly* the same.

There is precedent for thinking that higher-order beliefs and evidence at least sometimes defeat first-order reasons.[15] When we have evidence concerning the quality of our evidence, we thereby have more evidence that bears on the probability of the proposition in question. There are a few ways that evidence concerning the quality of one's evidence may impact the probability of a proposition for a subject. One way is this: whenever a person, considering whether p, obtains evidence *that there is* evidence that bears on the truth of p, that indirect evidence is evidence concerning p for that subject.[16] Anyone who thinks that how probable a proposition is depends upon the evidence for and against that proposition is thus committed to accepting that higher-order evidence bears on the *overall* probability of a proposition, and is not bracketed away as probability-neutral 'evidence.'[17]

Unless it is evidence that mysteriously has *no* bearing *at all* on the probability of the proposition in question, then higher-order concerns plainly *do* have an effect on the probability of the proposition for the subject. Surely we can, for theoretical purposes, distinguish different relationships that subjects' evidence may bear to propositions, and that is precisely what we see in the literature on higher-order evidence and defeaters. But bracketing for theoretical purposes does not imply commitment to the substantive claim that higher-order evidence and defeaters make absolutely *no* difference to epistemic probability.

Thus the Bandit's plausible claim to have distorted the subject's reasoning constitutes evidence for the subject that 2/3 is the wrong answer. The subject,

---

[15] David Alexander, "The Problem of Respecting Higher-Order Doubt," *Philosophers' Imprint* 13, 18 (2013): 1-12; Richard Feldman, "Respecting the Evidence," *Philosophical Perspectives* 19, 1 (2005): 95-119; Sophie Horowitz, "Epistemic Akrasia," *Nous* 48, 4 (2014): 718-744; Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, eds. Richard Feldman and Ted A. Warfield (New York: Oxford University Press, 2010), 111-174.

[16] Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers without God: Meditations on Atheism and the Secular Life*, ed. Louise Antony (New York: Oxford University Press, 2007), 194-214.

[17] Arguably, all evidentialists (broadly speaking) should think of epistemic probability as 'flat': evidence for p is evidence for p, whether that evidence comes in directly for p, or in the form of evidence about one's evidence for p, or evidence about one's evidence for one's evidence for p, and so on. Thanks to Lisa Miracchi for discussion on this point.

believing in accordance with the evidence, revises the probability of 2/3 downward from 1.0. Consequently, the subject's belief that 2/3 is the answer is lucky.

4. Inference and Probability

Steglich-Petersen's final argument that evidential certainty is insufficient for luck begins with this example from Martin Smith:[18]

> Background Color: Martin has set up his computer such that, whenever he turns it on, a random number generator determines the background color on his display. For one out of a million possible values, the background will be red. For the remaining 999,999, it will be blue. Martin turns on his computer, and then leaves the room before seeing the background color. A few minutes later, Martin's housemate Bruce enters the room and sees that the background color is blue.

Steglich-Petersen then applies the case to luck:

> Suppose that both Martin and Bruce, in light of their respective evidence, form the belief that the background is blue. And indeed it is blue. Are any of them lucky to have formed a true belief? Bruce's true belief does not seem lucky. He is looking at the display, and sees that the background is blue. If it hadn't been blue, he wouldn't have believed it. Very plausibly, Bruce thus knows that it is blue… Martin, on the other hand, seems to enjoy at least some degree of luck in his true belief, even if it was statistically very unlikely to be false. Martin would not be justified in believing outright that the screen is blue.[19] He does not know that it is blue, only that it is very likely to be blue. And if it hadn't been blue, he would still have believed it to be blue. So it does seem at least slightly lucky that he has ended up believing the truth. At the very least, the following seems clear: Martin is luckier than Bruce is, in ending up with a true belief.[20]

Steglich-Petersen's point is that knowledge requires more than evidential probability of 1.0. Bruce's belief has these additional knowledge-making properties, whatever they are, while Martin's belief is based merely on probability. So, Bruce knows and Martin does not.[21]

---

[18] Martin Smith, "What Else Justification Could Be," *Nous* 44, 1 (2010): 10-31; Martin Smith, *Between Probability and Certainty: What Justifies Belief* (New York: Oxford University Press, 2016).

[19] Steglich-Petersen continues after the passage: "As Smith notes, in cases like this, it seems that while Bruce would be justified in believing that the background color is blue, Martin would not. Martin would of course be justified in believing it to be *very likely* that the color is blue, but not justified in outright believing this" (p. 9). Note that this diagnosis is correct only if one is justified in believing only what one has *infallible* reasons to believe.

[20] Steglich-Petersen, "Knowledge or Certainty," 10.

[21] I am willing to entertain the example for the sake of argument, but it should not be forgotten that luck requires significance. It is rather implausible that the screen color is significant to

The same verdict can be explained probabilistically. It is stipulated in the case that Bruce *sees that* the screen is blue, while Martin infers the color from his knowledge of the color-generating algorithm. So, we can describe the conditional probability that the screen is blue for each subject this way:

For Bruce, Pr(screen is blue | Bruce sees that it is blue) = 1.0

For Martin, Pr(screen is blue | there is a 999,999/1,000,000 chance that it is blue) = 0.999999

Surely if one *sees that* p then the probability for that person that p is 1.0.[22] If that is correct, then this case provides no reason to think that KA explains what EAL cannot, because *both* views secure the intuitive verdict: Martin is luckier than Bruce.

To press the point, one must supply a reason as to why we should think that the probability of Bruce's belief is lower than Martin's. Steglich-Petersen suggests that Bruce is unable to eliminate various possibilities of deception: that he is hallucinating, that the lighting is tricky, etc., and that such uneliminated possibilities affect the probability of Bruce's belief. But surely these possibilities *are* eliminated if *seeing-that* is factive. Surely if Bruce sees that the screen is blue then he knows that the screen is blue, and if he knows that the screen is blue, then it is impossible relative to everything Bruce knows that the screen is *not* blue (or that there is no screen, or that he is dreaming, and so forth). The objection still fails because the probability of Bruce's belief is higher than Martin's.

5. Luck and Evidence

Even if the previous objection fails—the objection was that more than evidential probability matters for luck—the strength of Steglich-Petersen's argument forces me to admit that my description of the EAL as claiming that luck is determined by significance and a subject's evidence *alone* is slightly misleading, as there really are factors that contribute to justification (and knowledge) that are in a way different

either of these subjects. One might say with Pritchard (according to his updated view that luck does not require significance) that *risk* is present even if luck is not: Duncan Pritchard, "The Modal Account of Luck," *Metaphilosophy* 45, 4-5 (2014): 594-619. Also see Fernando Broncano-Berrocal, "Luck as Risk," in *The Routledge Handbook of the Philosophy and Psychology of Luck*, ed. Ian M. Church and Robert J. Hartman (New York: Routledge, 2019).

[22] In his defense of epistemological disjunctivism, Pritchard denies that seeing that p entails knowing that p. If that were correct (it is not), it would apply equally to Steglich-Petersen's diagnosis of this case and to the argument I just gave. Cf. Duncan Pritchard, *Epistemological Disjunctivism* (Oxford: Oxford University Press, 2012).

from just first-order *evidence*.[23] We can see this by comparing the role of additional factors in justification with the extra premise used to generate a vicious regress in "What the Tortoise Said to Achilles."[24] Achilles attempts to infer C from

1. If A and B, then C

2. A and B

using modus ponens. The inference is plainly valid, but the Tortoise challenges him to *prove* that the inference is legitimate by adding a new premise that states the modus ponens rule:

3. 'If [A and B, then C] and [A and B] then C'

But by granting the Tortoise's request, Achilles has conceded enough for the Tortoise to show that Achilles is committed to a vicious regress. The requirement that from any set of premises that entail a conclusion one *also* needs a *premise* that describes the *rule* that permits the inference begins an infinite series that prevents the conclusion from ever being legitimately inferred.

We could, using the language of 'evidence,' say that Achilles's evidence for C consisted of just two things: premises 1 and 2 of the argument. But then the Tortoise would rightly point out that Achilles has no business inferring C without at least some recognition of the *connection* between 1 and 2, and C. If awareness of that connection is counted as 'evidence' *in the same way* as belief in the premises 1 and 2, then the Tortoise's vicious regress begins. The premises and the awareness of the relation between premises and conclusion are two different things, and it is misleading to call both of them "evidence" as though there is no difference between them.

With this in mind, we should revisit Steglich-Petersen's argument that the Logical Bandit case cannot be shown to be lucky because the probability of the 2/3 answer is below 1.0 given the subject's evidence. (Let us also grant what I have argued above is impossible: that the subject understands the details of the case while being unable to do more than guess blindly at the answer.) We *could* categorize the factors that contribute to the subject's justification for believing that 2/3 is the answer this way:

The subject's evidence = the details of the puzzle

The subject's non-evidential justification-contributor = awareness that the

[23] I should have seen this problem coming, since I have used the same distinction to argue against a different thesis, in Bryan C. Appley and Gregory Stoutenburg, "Two New Objections to Explanationism," *Synthese* 191, 7 (2016): 1391-1407.

[24] Lewis Carroll, "What the Tortoise Said to Achilles," *Mind* IV, 14: 278-280.

evidence entails 2/3

If we describe the factors that contribute to the subject's justification in this way, then appeal to probability on the subject's evidence does not solve the Bandit case. We would then be conceding that the subject's awareness of the connection between the evidence and the proposition has no bearing on epistemic probability.

However, while there is an important difference between evidential and non-evidential justifiers, there is a much more substantial difference between contributors to a subject's justification that have to do with the *truth* of the proposition believed, and those that do not, like a subject's level of confidence or the stakes at play in a context. Indeed, I introduced talk of evidential probability only when distinguishing my epistemic probability account from the *interest-relative* version of Steglich-Petersen's view.[25] If we draw the distinction between what counts as evidence and what does not count as evidence by factors that contribute to the truth of a proposition and those that do not, then for the purpose of talking about how probable a proposition is for a subject, it is reasonable to collapse the evidential/non-evidential justifier distinction as it was drawn above. Reasonable, but still misleading.

So, if we use 'evidence' broadly to include all and only those factors that bear on the truth of a proposition, factors like a subject's experiential states, knowledge, justified beliefs, and awareness of evidential connections, we can say that on this *evidential* probability theory of luck the probability that the answer is 2/3 is indeed below 1.0, since the subject fails to appreciate the connection between those factors that support the 2/3 answer and that answer being correct (even granting, once again, that the subject understands the puzzle).[26] Since the subject's awareness of the connection between the answer and the other factors that support that answer is one of the factors that supports the truth of the 2/3 answer, the overall probability for the subject is below 1.0.

6. Epistemic Accounts of Luck and Epistemic Luck

A long-standing hope of the literature on luck is that an analysis of luck will provide insight into some lasting philosophical problems, like how we can have

[25] Stoutenburg, "In Defense," 6.

[26] Eventually it will be harder to be this inclusive about evidence, as philosophers disagree on what constitutes knowledge and justification. Safety theorists, for instance, have a modal condition on knowledge. So, if 'S knows that p' entails that p is safe, and p is a part of S's evidence, then a modal fact constitutes some of S's evidence. This is an example of how specific accounts have implications for more general ones. Philosophers interested in an epistemic probability account of luck can take sides on this issue as they wish.

moral responsibility when so much of what we are and do is lucky, and why we cannot have knowledge when the truth of a belief is due to luck.[27]An epistemic account of luck is at least initially appealing in this regard: we have the intuition that a belief that is true by luck is not knowledge, and then we add that by 'true by luck' we mean that the subject's belief suffers a distinctively *epistemic* failing.

The initial appearance is deceiving, however, if by "B is true by luck" we just mean "B is not knowledge," as KA would have it. Then the account is saying what we knew all along: that when a belief is not knowledge, it is not knowledge. This trivial equivalence is all the insight that a knowledge account of luck can give us. Whatever other virtues such an account could have, this consequence is disappointing, and prevents KA from having much theoretical significance.[28]

In contrast, an epistemic account of luck understood in terms of epistemic probability can identify an interesting and principled connection between luck and knowledge. We can say that a belief is true by luck when the epistemic probability of the belief falls below a specific threshold, most plausibly (in my view) the limit of epistemic probability: 1.0. Any significant true belief is at least *very slightly* lucky if the probability of the belief on the subject's evidence is below 1.0.

This way of thinking about *epistemic* luck illuminates some connections between epistemic probability and knowledge. If the view is true, then underdetermination arguments for skepticism, lottery problems, closure-based skeptical arguments, and Gettier scenarios all have in common that in each, the truth of a subject's belief is to a slight degree due to luck, because the subject's basis for belief does not guarantee the truth of the proposition believed. It is significant that an epistemic probability theory of luck can unite a number of interestingly similar issues about knowledge.

One might worry that this account of epistemic luck immediately leads to skepticism. Here is one defective argument for that conclusion:

1. S cannot know that p if the truth of p is just a matter of luck.

2. The truth of p is just a matter of luck when $P(p) < 1.0$.

3. So, we cannot know that p.

For the argument to be at all forceful would also require the stipulation that most ordinary beliefs are like p.

---

[27] See Ballantyne for arguments that this goal of the luck project is misplaced: Nathan Ballantyne, "Does Luck Have a Place in Epistemology?" *Synthese* 191, 7 (2014): 1391-1407.

[28] Steglich-Petersen is aware of this kind of limitation of his view ("Luck," 376). Here, I am arguing that the limitation has implications for the significance of the view.

The argument is defective because premise 2 is false. It is highly misleading to say that the probability of p is "just a matter of luck" when p is *highly* probable but still below 1.0. To say an event is "just a matter of luck" implies that the occurrence of the event is due *only* to luck and to nothing else. But that is not what EAL says, even if combined with the idea that a belief is true by luck when it is not guaranteed by the subject's evidence: it may be that the subject's belief is 0.99 probable, and true. That would not make the truth of the belief "just a matter of luck." It would make the belief .01 a matter of luck, and .99 a matter of evidential support.

Premise 1 is not implied by the epistemic probability account of luck. Nothing in that account of luck implies that we must accept that knowledge is incompatible with even the slightest trace of luck. Even those who think that luck and knowledge are incompatible, broadly speaking, should still clarify *how much* luck is tolerable. Faillibilists might accept the epistemic probability account of luck and say that knowledge is compatible with some small degree of luck, and infallibilists can deny this. Either way, the truth of premise 1 is required for the argument to be successful, and nothing about the epistemic probability theory of luck itself implies that premise 1 is true. But if one thinks that absolutely no luck is compatible with knowledge, then the skeptical conclusion follows.

7. Conclusion

An epistemic probability account of luck successfully explains how beliefs about necessary truths can be lucky and delivers the correct verdict about luck when considering the epistemic status of beliefs formed by vision and through inference. While arguing for these claims in the course of responding to objections from Steglich-Petersen, I also further developed my epistemic probability theory of luck. I argued that higher-order evidence affects epistemic probability, clarified the conditions under which an experience is taken up as evidence, distinguished ways that factors relevant to justification affect epistemic probability, argued that the epistemic analysis of luck does not by itself imply skepticism, and underscored the impressive explanatory power of the epistemic analysis of luck. EAL remains a fruitful and viable option for a unified theory of luck.