Is Lucky Belief Justified?

Fernando Broncano-Berrocal
(University of Barcelona – LOGOS)

Abstract: The main lesson from Gettier cases is that while one cannot know a proposition by luck, one can hold a lucky true belief justifiedly. Possibly because the latter is taken for granted, the relationship between epistemic justification and epistemic luck has been less discussed. The paper investigates whether luck can undermine doxastic justification, and if so, how and to what extent. It is argued that, as in the case of knowledge, beliefs can fall short of justification due to luck. Moreover, it is argued that justification-undermining luck is a problem for both internalist and externalist conceptions of justification. Accordingly, it is shown that epistemic luck is a more widespread phenomenon than many in epistemology commonly assume.

Orthodoxy in epistemology has it that one *cannot* know a proposition by luck, but also that one *can* hold a lucky true belief *justifiedly*. Gettier cases prove both points (Gettier 1963). When Gettierized agents are in the process of forming false but well-justified beliefs, accidental factors make it the case that they end up believing true propositions. Their beliefs are not knowledge—because they are true by luck—, but they are still justified. Possibly because this is something that epistemologists usually take for granted—namely, that justification can be preserved when luck undermines knowledge—less discussion has been devoted to the relationship between epistemic justification and epistemic luck.

In this paper, I aim to investigate whether luck can undermine epistemic justification, and if so, how and to what extent. In particular, I will explore how doxastic justification—as conceived by some representative internalist and externalist theories—can be undermined by luck. This kind of project may be particularly surprising in the case of internalist justification. One reason is that internalists are primarily concerned with justification rather than knowledge, and it is typically assumed that epistemic luck undermines knowledge, not justification. As I will argue, however, this generalized assumption is wrong: epistemic luck can undermine doxastic justification when the notion is interpreted in both externalist and internalist ways.

The plan is the following. In §1, I survey the current state of the art on epistemic luck from a critical methodological perspective. In §1.1, I examine how different types of knowledge-related luck—henceforth *k-luck*—are distinguished. In §1.2, I argue that while k-luck is typically

understood in modal terms, its existence does not depend on the correctness of any account of luck or epistemic luck. In other words, we should not be persuaded to give up our intuitions about the presence of luck or epistemic luck in a given case just because a theory predicts that there should be none. In §1.3, following the same kind of standard methodology as in the case of k-luck, I consider analogous varieties of justification-related luck—henceforth *j-luck*. In §1.4, I discuss whether there are truth-related variants of j-luck. In §1.5, I argue that we should not worry that the proposed way of individuating forms of j-luck makes the notion too theory-dependent to be of philosophical interest. In §2, I argue that j-luck can undermine internalist justification, and in §3, externalist justification. The views of epistemic justification that I draw on for illustration are accessibilism and evidentialism (in the case of internalism) and process and virtue reliabilism (in the case of externalism). In addition, I examine how these views can eliminate dangerous j-luck. This provides a general idea of what is needed to develop an anti-j-luck epistemology. In §4, I make some concluding remarks.

1 Types of k-luck and j-luck

1.1 A methodological survey of the literature: how types of k-luck are distinguished

In attempting to motivate his account of knowledge as *non-accidental belief*, Unger (1968) distinguishes several accidents concerning factors that are typically thought to enable or cause the acquisition of propositional knowledge. Following Unger, Pritchard (2005) translates these accidents into a taxonomy of types of epistemic luck—or, as Unger would put it, 'epistemic accidents'—, which shows, among other things, that not all types of epistemic luck are incompatible with knowledge.[1]

Unger and Pritchard's methodological approach to distinguishing between dangerous and harmless varieties of k-luck is based on intuitions and is fairly common in epistemology. Briefly, the methodology is as follows: for any factor that is thought to be necessary for knowledge (e.g., truth, evidence, and so on), the type of k-luck corresponding to that factor is considered epistemically harmless (compatible with knowledge) if in all cases where it obtains by luck knowledge is not undermined; in contrast, if knowledge is lost, it is considered epistemically dangerous (incompatible with knowledge).[2]

---

[1] See also Engel (1992) and Vahid (2001) for discussion of types of k-luck.
[2] This methodology can be challenged in a number of ways, such as by doubting the idea that conceivability is a guide to possibility, by criticizing the method of cases in epistemology, or by questioning the reliability of philosophical intuitions

The following is a non-exhaustive list of factors typically thought to enable or determine the acquisition of knowledge—and which may obtain by luck. S knows that p only if:

(K1) S exists.

(K2) S possesses the relevant physical and psychological constitution.

(K3) S forms the belief that p.

(K4) There is a truth-maker for p / the proposition p is true.

(K5) S comes to believe p truly.[3]

(K6) S has evidence E that supports p.[4]

Factors K1 and K2 are affected by what Pritchard calls:

Capacity luck: It is a matter of luck that S is capable of knowledge.

The formulation of capacity luck is meant to cover both the luck that an agent has to exist (i.e., luck concerning K1) and the luck that an agent has to possess the physical and psychological constitution required for knowledge (i.e., luck concerning K2). Other types of epistemic luck are:

Doxastic luck: It is a matter of luck that S believes p (K3).

Content luck: It is a matter of luck that there is a truth-maker for p, i.e., it is by luck that p is true (K4).

---

in general. Addressing these objections is beyond the scope of this paper. An alternative methodology is to test folk intuitions about widely discussed cases in epistemology and theorize about the results (e.g., Turri et al. 2015). Of course, this assumes that people are able to reliably individuate intuitions about luck, as opposed to, say, risk, chance, randomness, or similar concepts. Whether people are so reliable is largely an open question.

[3] (K3), (K4), and (K5) are different necessary factors for knowing a proposition p. For (K3), what matters is whether the agent adopts a doxastic attitude with content p, i.e., regardless of whether that content is true or false. For (K4), what matters is whether there is a truth-maker for p, i.e., regardless of whether there is, in addition, a doxastic state that has p as content. (K5) is the fact that the agent comes to believe p truly. Preconditions for (K5) are, of course, (K3) and (K4). In other words, (K5) can hold only if (K3) and (K4) hold. And if (K3) and (K4) hold, (K5) also holds. In other words, S comes to believe p truly if and only if S forms the belief that p and there is a truth-maker for p. Given this equivalence, it might be redundant to introduce (K5) as a distinctive necessary factor for possessing propositional knowledge. However, it is useful to do so because it facilitates the identification of a special kind of knowledge-undermining luck, namely veritic luck.

[4] Insofar as evidence is typically considered a mental-dependent (internalist) notion, hard-core externalists will reject that posession of evidence is necessary for knowledge. However, since the point here is not to define knowledge but to distinguish interesting varieties of epistemic luck, it is useful to consider factors *typically* associated with knowledge.

> *Veritic luck*: It is a matter of luck that S comes to believe p truly (K5).

> *Evidential luck*: It is a matter of luck that S obtains evidence E that supports p. (K6).

Pritchard, following Unger's discussion, argues that most kinds of k-luck are epistemically harmless: only veritic luck is truly incompatible with knowledge. To see this, consider the following series of cases, all of which involve knowledge:

> (A) At the beginning of time, God flips two coins: first to decide whether to create Adam or Peter; then, to determine whether the created being should have eyes to see the world. As luck would have it, God creates Adam and not Peter, and fortunately for him, he is endowed with the ability to see. He opens his eyes and comes to believe, correctly, that he is surrounded by apple trees.

(A) is a case of capacity luck, in that it is by luck that Adam is capable of knowledge, both in the sense that it is a matter of luck that he comes to exist (K1) and that he is endowed with a reliable visual system (K2).

> (B) The next coin God flips is to decide whether to make roses red or blue. Roses become red. Adam, who has never seen a rose in his short life, comes across a rose bush and believes, correctly, that there are red roses in Paradise.

(B) is a case of content luck: the state of affairs that makes Adam's believed proposition true (that roses are red) obtains by sheer luck (after a coin toss).

> (C) Time goes by. Someone is robbing the Bank of Paradise. As the robber flees from the bank, his mask slips off for a few seconds allowing Adam to see that it is Cain who has just robbed the bank.[5]

(C) is both a case of evidential and doxastic luck. First, it is a stroke of luck that there is available evidence in support of the proposition that Cain has just robbed the bank (it could easily have been the case that Cain's mask had not slipped off), but precisely because Adam comes across lucky evidence and he takes it at face value, it is also by luck that he ends up with a belief about the identity of the robber.

---

[5] See Nozick (1981: 93) for the original case.

(A), (B) and (C) are intuitive cases of knowledge, i.e., capacity, content, evidential, and doxastic luck are compatible with its possession.[6] But compare:

> (D) When strolling around Paradise, Adam sees what looks like a snake in the tree in front of him and becomes convinced that there is a snake in the tree. What Adam saw was actually a snake-looking branch. Just before he formed his belief, however, a snake had accidentally fallen from a higher tree and bounced against several branches until it finally got caught in the lower branches on the unseen side of the trunk. Adam ends up forming a true belief, but could easily have been false.[7]

(D) is a paradigmatic case of veritic luck, because it is a matter of luck that Adam comes to truly believe that there is a snake in the tree. However, the fact that (D), a case of ignorance, is also a case of veritic luck is not yet sufficient to establish the general claim that all instances of veritic luck are cases of ignorance, but only that veritic luck, though perhaps not incompatible with knowledge, is epistemically dangerous. However, everything points to their incompatibility. After all, it is surprisingly difficult to find clear cases of veritic luck that are not cases of ignorance. If conceivable, the resulting cases are too convoluted to generate stable intuitions about the presence or absence of knowledge.

The incompatibility of knowledge and veritic luck is also underpinned by a historical fact: after Gettier's seminal paper, many diagnoses of why Gettierized subjects lack knowledge appeared; the justification condition was accordingly tweaked (or replaced) in all sorts of ways to solve the problem. But however successful or unsuccessful these solutions may have been, the central idea that has stayed put since then—to the point that it is deeply entrenched in current epistemological theorizing—is that one cannot have knowledge if one comes to believe a true proposition by luck.[8]

---

[6] The case of doxastic luck is somewhat more complicated. We have seen that there are cases of doxastic luck and knowledge, so that doxastic luck is compatible with knowledge. However, there are also cases of doxastic luck and ignorance. Suppose that God hypnotizes Adam and tosses a coin to decide whether to induce in him the belief that p or the belief that not-p, where p is true. As luck would have it, Adam ends up believing p when he wakes up. While this is a clear case of doxastic luck, it is unclear whether Adam has knowledge. This, however, does not necessarily mean that doxastic luck is epistemically dangerous, because the factor that undermines knowledge does not seem to be luck, but the fact that the belief in question is not formed by the exercise of a cognitive faculty. In case (C), on the other hand, the target belief is formed in an epistemically appropriate way, namely by following the available evidence. This is plausibly why it is a case of knowledge.

[7] The case is analogous in all relevant aspects to Chisholm's well-known sheep-in-the-field case (Chisholm 1977: 105).

[8] See Baumann (2012) for a dissenting opinion.

1.2 General views of luck are not relevant for distinguishing types of k-luck

Intuitions about how luck undermines knowledge have been present in epistemology long before accounts of luck and epistemic luck were developed. Just as the existence of k-luck does not depend on the correctness of a proposed solution to the Gettier problem, neither does it depend on the correctness of such views.

The different views of luck and epistemic luck on the market (e.g., modal, probabilistic, lack-of-control views) provide necessary and sufficient conditions for luck/epistemic luck, which means that the proposed definitions make predictions about the application of the concepts of luck and epistemic luck to particular cases. If we have the clear intuition in a particular case that the target event—e.g., winning the lottery or forming a true belief—is due to luck and, at the same time some of the necessary conditions of a particular view are not met, we have reason to believe that such a view does not capture our intuitions about the presence of luck or epistemic luck. If, on the other hand, we have a clear intuition that the target event does not result from luck, and yet all the conditions of a particular view are satisfied, we have reason to believe that this view does not capture all it takes for an event to be epistemically or non-epistemically lucky. Accordingly, one thing should be clear about this kind of methodology: intuitions about luck come first. Thus, if the prediction of an account of luck or epistemic luck does not match our intuitions about the cases, our intuitions are not wrong: the account is.[9]

For illustration, consider the prevailing view of epistemic luck, the *modal view*, which says that a belief that p is true by luck—i.e., veritically lucky—just in case in most close possible worlds in which one forms a belief in p in the same way as in the actual world, p is false.[10] Consider now a version of case (D), a paradigmatic Gettier-style case:

---

[9] At this point, theorists of luck whose accounts are inconsistent with folk intuitions can either introduce an error theory to explain why we mistakenly apply or do not apply the concept of luck when we should not, or they can opt for an empirical investigation of folk intuitions and try to derive some general conclusions about how luck actually works to support their theoretical claims—see e.g., Turri et al. (2015).

[10] See Pritchard (2005) for a seminal defense of this view. The modal account of epistemic luck derives from the more general modal account of luck—see Pritchard (2014) for a recent defense—, which holds that an event is lucky just in case it occurs in the actual world, but would not occur in most nearby possible worlds in which the relevant initial conditions for its occurrence are the same as in the actual world.

(D*) Everything is as in case (D), except for the fact that the snake has always lived on the branches of the unseen side of the tree (and has never left it and never will).

It is still by luck that Adam gets things right. After all, it is a *coincidence*, and therefore luck, that (i) Adam sees a snake-looking branch and thus becomes convinced that there is a snake in the tree while (ii) there is actually a snake on the unseen side of the tree trunk. But the snake has always been there, and will always be there. There is not the slightest chance that Adam would have been mistaken. In other words: in *no* close possible world, Adam's belief would be false.

In contrast, much argument would be needed to refute the intuition that (D*) is a case of knowledge-undermining luck. First, (D*) is a standard Gettier-style case, and Gettier-style cases are paradigmatic examples of dangerous epistemic luck. Moreover, as noted above, the case is a clear case of luck out of coincidence, where a coincidence is composed of events that are produced by independent causal factors in a way that we cannot explain why they come together (Owens 1992; for discussion, see Lando 2017). In (D*), (i) the fact that Adam forms his belief by looking at a snake-looking branch and (ii) the fact that a snake lives in the tree are not connected in this sense. However, it does not matter that the components of a coincidence are modally robust—so that both would occur in most nearby possible worlds—: this does not make it any less of a coincidence. All coincidences, however, involve luck.

The upshot is that we should not be persuaded to abandon our intuitions about the presence of luck in a particular case just because a theory (in this case the modal view, but the same point applies to any view) predicts that there should be no luck involved. Intuitions about luck are not defeated in this way.[11] The same is true for epistemic luck. The difference between ordinary and epistemic luck, after all, is simply that the latter affects factors that are necessary for knowledge or justification. In other words, there is nothing special about epistemic luck except the fact that the relevant lucky factors are epistemic.

With the preceding methodological considerations in mind, let us now analyze the factors typically considered necessary for justification in order to distinguish different types of j-luck.

---

[11] Perhaps more arguments are needed to prove that the modal view is wrong. For some of its predictions are consistent with our intuitions about luck, namely cases of modally fragile lucky events, i.e., actual events that could easily have failed to materialize, such as winning the lottery.

1.3 Analogous types of j-luck

Propositional justification (i.e., having justification for believing p) does not require believing in p, so luck affecting belief formation is not relevant to this type of justification. In contrast, doxastic justification (i.e., justifiedly believing p) requires belief in p. In what follows, I will focus only on lucky factors that influence doxastic rather than (only) propositional justification. The reason for this is that the influence of luck on propositional justification has already been studied in the literature. For example, Bondy and Pritchard (2017) argue that epistemic luck, and in particular what they call *propositional luck*, can affect propositional justification for p in ways that affect the adequacy of the epistemic basis of the belief that p (in general). More specifically, they argue that individuals who have true beliefs for which they have good evidence and against which there is no counterevidence, but who do not believe on the basis of the available good evidence, are lucky to have propositionally justified beliefs (i.e., beliefs for which they have good evidence).

My aim is to take the discussion of luck in the context of epistemic justification a step further, arguing that the *way* we form beliefs may be subject to luck in a way that precludes justification even if those beliefs are properly based. For example, if a subject forms a true belief by following a correct rule of inference, the following of the correct rule may still involve too much luck for the subject to be justified in her belief. Moreover, I will argue that the resulting form of j-luck in such cases of proper basing is capable of undermining doxastic justification as conceived by prominent internalist and externalist views.[12]

Let us first consider some necessary conditions that doxastic justification has in common with knowledge. S's belief that p is justified only if:

>   (J1) S exists.

>   (J2) S possesses the relevant physical and psychological constitution.

>   (J3) S forms the belief that p.

And although the following condition is not uncontroversial (see a more specific version in §2.3):

>   (J4) S has evidence E that supports p.

---

[12] For the converse project of characterizing the distinction between internalism and externalism of epistemic justification in terms of varieties of j-luck, see de Grefte (2018).

As in the case of k-luck, (J1) and (J2) allow to distinguish a distinctive type of j-luck:

> *Capacity j-luck*: It is a matter of luck that S is capable of justified belief.

We have seen that capacity luck is compatible with knowledge possession. Nor does capacity j-luck undermine doxastic justification. In both cases, even if one is properly constituted physically or psychologically by luck, this initial luck is not inherited by the resulting epistemic standings—knowledge or justified belief. In general, it does not necessarily follow from the fact that it is by luck that one is capable of φ-ing, that one φ-es by luck; and this is also true of the epistemic case.[13]

Consider now (J3). Propositional justification does not require one to believe in p, so no luck affecting belief formation is relevant to this kind of justification. In contrast, doxastic justification requires the formation of a belief in p, which means that we can distinguish a variety of j-luck that influences this factor:

> *Doxastic j-luck*: It is a matter of luck that S believes p.

Is doxastic j-luck compatible with doxastic justification? We have seen that cases of doxastic luck and lack of knowledge are plausibly not due to the luck involved, but to epistemically inappropriate factors that affect the way the relevant beliefs are formed (see fn. 6). The same is true in the case of justified belief. If you come to believe that p after being hypnotized by a hypnotist who flipped a coin to decide whether to induce in you the belief that p or the belief that not-p, what prevents you from justifiably believing p is not the luck involved but the fact that you were led to form a belief by hypnosis.

Finally, consider (J4). Evidential luck in the case of knowledge is the same as evidential luck in the case of justified belief as long as this kind of epistemic luck is understood as luck affecting the *acquisition* of evidence:

> *Evidential j-luck*: It is a matter of luck that S obtains evidence E that supports p.

Cases like (C) motivate the idea that this kind of j-luck is epistemically harmless with respect to the formation of justified beliefs. Of course, the acquisition of evidence is not the only relevant factor in (J4): as we will see in §2.3, *evidential support* (as well as basing on the evidence) can also

---

[13] See Coffman (2009) for arguments for the falsity of this thesis, which he calls the 'luck infection' thesis.

come about through luck, and in this case the relevant variant of evidential j-luck is not so epistemically harmless.

1.4 No truth-related varieties of j-luck?

Many epistemologists hold that epistemic justification is required for knowledge, but few that knowledge is required for epistemic justification. For example, knowledge-first reductionists hold that one justifiably believes that p if and only if one knows that p (Sutton 2007). Others are knowledge-firsters but not reductionists. For example, Williamson (2000) holds that justified beliefs can only be justified by known propositions, which does not necessarily imply that justified beliefs are known beliefs. *Pace* knowledge-first reductionism, we should not expect doxastic justification to have exactly the same necessary conditions as knowledge. Since only factors necessary for the epistemic standings in question can be used to distinguish epistemologically interesting forms of epistemic luck, we should also not expect knowledge and justification to be affected by the same kinds of epistemic luck. In particular, whether there are truth-related varieties of j-luck, such as content or veritic luck in the case of knowledge, depends on whether justification is factive, a claim that knowledge-first reductionism entails. Proponents of such a view face two nontrivial tasks, however: they must argue against the consensus in epistemology that Gettier-style cases are cases of justified belief without knowledge (Sutton 2007 makes such an attempt) and against the core intuition of internalism that subjects in evil demon worlds can have justified beliefs (see below). For the purposes of this article, I will stay on the orthodox side and assume that justification is not factive and that there are therefore no truth-related varieties of j-luck. But none of the arguments I will offer depend on this assumption.

1.5 Concluding methodological remarks: no troubling theory-dependence

The next step in distinguishing further variants of j-luck is to consider necessary conditions as proposed by canonical views of epistemic justification. If a view says that S's belief that p is justified only if condition C is satisfied, then C can be used to distinguish a particular form of j-luck. Following the same methodology as in the literature on k-luck, the idea is to see what happens to our intuitions about the presence of justification when C occurs by luck. If justification is undermined, we have reason to believe that the resulting form of j-luck is epistemically dangerous. If it is not—and we cannot imagine a plausible case in which C is satisfied while epistemic justification is absent—we can confidently regard it as epistemically harmless.

One might worry that this way of distinguishing forms of j-luck makes the notion too theory-dependent to be of philosophical interest. Note, however, that the methodology is the same as in the case of k-luck: all varieties of k-luck are theory-dependent in exactly the same sense. After all, their existence depends on certain factors that are necessary for knowledge. For example, veritic luck is assumed to exist because knowledge is factive. Of course, almost no one in the business of analyzing knowledge (nor any knowledge-firster) would claim that knowledge is not factive.[14] But precisely because there is no view that presupposes this, veritic luck automatically becomes a problem for everyone: it arises from a necessary condition that everyone accepts.

The only difference with the knowledge debate is that the conditions on justification that help distinguish dangerous forms of j-luck are somewhat more controversial, since they depend to some extent on irreconcilable epistemic sensibilities (internalism, externalism). This does not mean, however, that the resulting forms of dangerous j-luck are theory-dependent any more than their knowledge-related cousins are: they all involve theoretical commitments.

2 J-luck and internalism

2.1 GOOD and BAD

To see how epistemic luck can undermine internalist justification, let us first consider perhaps the most important intuition motivating internalism in general, namely the idea that there is no difference in justification between someone who is, say, a good reasoner under normal circumstances and someone who is a good reasoner in evil demon worlds, where the assumption is that in such worlds all beliefs are false and hence unreliably formed. Interestingly, internalist orthodoxy also assumes that in evil demon worlds some beliefs are *not justified*. Stewart Cohen here:

> We can imagine two inhabitants of this [evil demon] world, *A*, who is a good reasoner, i.e., reasons in accordance with the canons of inductive inference, and *B*, who engages in confused reasoning, wishful thinking, reliance on emotional attachments, guesswork, etc. (...) *A*'s beliefs are conditioned by the evidence whereas *B*'s beliefs are not. *A* is a good reasoner whereas *B* is not. *A*'s beliefs are reasonable whereas *B*'s beliefs are not. There is a fundamental epistemic difference between the beliefs of *A* and the beliefs of *B* (...) [which] is marked precisely by the concept of justified belief. Beliefs produced by good reasoning are paradigm cases of justified

---

[14] But see Hazlett (2010).

belief and beliefs arrived at through fallacious or arbitrary reasoning are paradigm cases of unjustified belief. Whether or not reasoning results in false belief, even if this happens more often than not, is irrelevant to the question of whether the reasoning is good. (Cohen 1984: 283)[15]

A and B in Cohen's quote may well be one and the same person: it is possible for a person to be a good and a bad reasoner, at least as long as both things are not the case at the same time. Maybe you made a lot of logical mistakes in high school, but then studied philosophy, got a PhD in logic, and started making valid inferences. Are your beliefs in the conclusions of your inferences justified after you earn your PhD? Plausibly, the answer is yes. You were a poor thinker, but now you have improved your reasoning skills. To deny that your beliefs are justified contradicts the idea that learning improves your epistemic position. This does not mean that your epistemic position cannot deteriorate at a later time: you might, for example, start reasoning badly again because of dementia. Nevertheless, your beliefs in the conclusions you reached after your PhD and before your dementia are justified. I think this claim is uncontroversial. And we can express it in general internalist terms as follows:

GOOD: S's belief that p is justified at t if (i) S does everything right according to internalist standards in forming the belief that p at t and (ii) S would or could *not* easily have failed to satisfy these standards when settling the question of whether p under the same kind of circumstances at earlier or later times in the vicinity of t, even if S would or could have failed to meet those standards at more distant times or under relevantly different circumstances.

What if the intervals in which one meets the relevant internalist standards are not so long that it would be very easy not to meet them? In that case, justification is plausibly undermined. I think that if we accept GOOD, we should also accept its negative counterpart:

BAD: S's belief that p is *not* justified at t if (i) S does everything right according to internalist standards in forming the belief that p at t and (ii) S would or could easily have failed to satisfy these standards when settling the question of whether p under the same kind of circumstances at earlier or later times in the vicinity of t.

---

[15] The quote is from Madison (2010), who provides an excellent overview of internalist theories of justification.

Below, I will give an example to better illustrate GOOD and BAD, but first let us unpack these principles to avoid two possible misunderstandings.

1. Overall, GOOD and BAD capture the idea that justification, even when understood internalistically, requires some stability in the way one forms one's beliefs. If the methods by which one forms one's beliefs are suboptimal, e.g., if one is a poor reasoner, the beliefs one holds about *various propositions* will generally not be justified. However, even if this claim is generally true, there may be exceptions, or at least reasons to question it. Consider the case of a mathematics student taking a 200-question mathematics exam. Suppose this student screws up all but one question, justifying her answer to that question with the correct reasons and methods. Are we to conclude that just because the student messed up on the other questions, her belief in the correct answer is not justified? One might disagree: a good answer should be taken for what it is, namely a good (justified) answer to the question. GOOD and BAD, however, refer to what a subject believes or would believe about a single proposition. Therefore, both GOOD and BAD can in principle apply to the student and her good answer—we would need to fill in the details of the case further to know whether GOOD or BAD applies—*regardless* of her performance on the other questions (i.e., in relation to other propositions).

2. For GOOD or BAD to hold, the circumstances under which S would or could easily have failed to meet the relevant epistemic standards in settling the question of whether p must be the *same kind of circumstance*s under which S forms the belief that p at t. This proviso is meant to exclude cases in which S is by luck in a position to satisfy the relevant epistemic standards without necessarily satisfying them by luck (recall: it does not necessarily follow from the fact that it is by luck that one is capable of ϕ-ing, that one ϕ-es by luck). Suppose someone could have put something in your coffee that would have affected your ability to reason. But nothing of the sort happened. You were lucky that you ended up not being intoxicated, and therefore lucky enough to be in a position to deduce p correctly (you are a good reasoner, after all). Your belief in p is thus justified, and BAD says nothing against this fact precisely because, although you could easily have failed to meet the relevant epistemic standards because of the judgment-distorting substance, the circumstances under which you would have done so (i.e., while intoxicated) are different from the circumstances under which you would have met those standards (in the absence of intoxication). GOOD is, after all, applicable to you.

With these two caveats in mind, let us now consider the promised example that will hopefully make GOOD and BAD more intuitive, with Adam as the main character.

Adam lives in the evil demon world. Until college, he was always an ideological person and very poor at reasoning. In particular, when given evidence for the proposition p and, say, for the proposition that p entails q, he did not reason according to modus ponens, but believed q only when it fit his ideological values. Regardless of whether this kind of motivated reasoning is reliable (e.g., because the evil demon might change the world to match Adam's biased conclusions), motivated reasoning is the kind of reasoning-pattern that, by internalist standards, always leads to unjustified beliefs—just like "confused reasoning, wishful thinking, reliance on emotional attachments, guesswork, etc." Cohen (*ibid*.). And as far as Adam can tell, he argues well: he has no idea that he endorses the conclusions of his putative inferences only because he likes them or because they confirm or agree with his previous beliefs.

As the story progresses, however, Adam goes to college and takes a few courses in psychology and philosophy. As a result, he understands what motivated reasoning and modus ponens are. From that point on, Adam is careful to always reason according to the correct rules of inference so that his beliefs become justified (even though he is in an evil demon world and his beliefs are systematically guaranteed to be false).

Adam leaves college and applies what he has learned to his daily life. Whenever he finds evidence for p and p entails q, he makes sure to reason according to modus ponens rather than endorsing the conclusion just because q pleases him or because it agrees with his values. But the evil demon is truly evil and decides to create Twitter. Fooled by the hype, Adam becomes an avid Twitter user and increasingly gets caught up in echo chambers where only information that leans in a certain direction is discussed, and where little value is placed on arguing according to the proper rules of inference or on critical thinking in general. The result is this: sometimes Adam adheres to correct rules of inference, but sometimes he engages in motivated reasoning. Moreover, when he reasons according to modus ponens or some other rule of inference, he correctly thinks that he is reasoning according to that rule. However, when he engages in motivated reasoning, he likewise thinks, this time incorrectly, that he is following the appropriate rules of inference.

If the time interval in which Adam does everything right according to internalist standards (e.g., applies modus ponens to p and p entails q) is

sufficiently short, then the belief Adam forms in q is not justified if as a result he would or could very easily have failed to satisfy these standards (e.g., by engaging in motivated reasoning), which is BAD's prediction.

Although the shortness of a time interval may not in itself be sufficient to generate such an easy possibility of cognitive failure, it can arise when sufficiently short intervals of good reasoning *alternate* with sufficiently short intervals of bad reasoning. To see this, choose a long enough time interval so that GOOD applies (e.g., one, five, ten years). A long enough time interval is one in which the fact that one has formed unjustified beliefs in the past, or the possibility that one will form unjustified beliefs in the future, does not undermine the justification one has during that time interval. Then decrease the interval to a point where you clearly no longer have the intuition that beliefs are justified on the same basis (or under the same circumstances). The exact length is irrelevant for the purposes of the argument, since different people will have different intuitions about how long this 'short' interval should be. For example, some will think that it is sufficient for Adam to form beliefs that are internalistically okay one day, and to engage in unnoticed motivated reasoning the next day to view the earlier beliefs as unjustified. For others, however, the intervals may need to be shorter to intuit that in the 'good' intervals in which Adam plays by the internalist book, his beliefs are no longer justified. These good intervals could be, for example, a few hours long, but if need be, they can also be the shortest psychologically possible intervals (certainly less than a few hours).[16] Regardless of the length one intuitively considers appropriate,

---

[16] Empirical research on politically motivated reasoning (e.g., Calvillo et al. 2019; Gampa et al. 2019; Aspernäs et al. 2022) shows the influence of prior political beliefs on syllogistic reasoning, namely that individuals are more willing to accept conclusions that are consistent with their political beliefs than conclusions that are inconsistent (also known as belief bias in psychology). In some of these studies, participants had undergone logical training (Gampa et al. 2019; Aspernäs et al. 2022), so they understood how to reason well, and in one (Aspernäs et al. 2022) they were asked to try not to think too long before responding. This gives reason to believe that Adam's case is not unrealistic after all, and that it is not inconceivable that subjects may be prone to reasoning errors in short intervals, while they think they are reasoning correctly in the intervals in which they actually do so, but also in the intervals in which they manifest belief bias, of which they are unaware (because it is a bias)—further research also shows that rapid responding increases belief bias, which gives further reason to think that short-interval cases are not unrealistic; for example, in a study by Evans & Curtis-Holmes (2007) responding rapidly within 10 seconds, while did not inhibit all analytic processing of syllogisms, increased the amount of belief bias in participants.

beliefs based on good reasoning are not justified if they would or could *easily* have been formed by bad reasoning.

Thus, what leads to Adam's beliefs no longer being justified—and here comes a general diagnosis applicable to other possible cases as well—is not the brevity of the intervals per se, but the underlying fact that, in a scenario with such short intervals, it is completely *arbitrary* whether subjects comply with the relevant internalist norms. This arbitrariness arises from two facts in particular: (i) it is just as easy for subjects to form a belief based on biased reasoning as it is for them to do so by following correct rules of inference; (ii) they are unaware that they are using biased reasoning when they do so because they still think they are following the correct inference rules. This leads to a distinctive form of epistemic luck, namely *justification-undermining luck*.

Note the parallel with veritic luck in the case of knowledge: when agents are affected by veritic luck, it is just as easy for them to form true beliefs as false beliefs in the same propositions, and when the latter happens, they are *unaware* of it. The difference with the current case is that it is a matter of luck that Adam adheres to the appropriate internalist norms, not that he forms a true belief. In fact, this kind of luck is consistent with both Adam's beliefs being true or false, since Adam is an inhabitant of the evil demon world.

Another interesting parallel with veritic luck is the following. Veritic luck comes in two forms: *intervening* and *environmental luck* (cf. Pritchard 2012). Intervening luck occurs when one's belief in a true proposition is due to the intervention of luck rather than the exercise of one's cognitive faculties (case D would be an example). Environmental luck, on the other hand, occurs when one comes to believe a true proposition due to the exercise of one's cognitive faculties, but this could easily have not been the case (the fake barn case is the canonical example; cf. Goldman 1976). The kind of epistemically dangerous j-luck envisaged here is more akin to environmental luck than to intervening luck. For, from an internalist perspective, the subjects in question are doing everything right in their current situation, but they would or could easily have failed to do so.

To learn more about the nature of this kind of j-luck, we obviously need to know more about the specific internalist conditions for justification that are met when this kind of luck is present. This will tell us more about the theoretical adequacy of the various internalist views of justification that rely on such conditions, and the extent to which they may preclude this kind of dangerous j-luck.

Before analyzing specific internalist views, however, I want to anticipate the most obvious internalist response to this form of luck: the claim that in cases where BAD holds, the beliefs in question are justified precisely because the subjects in those cases comply with the relevant internalist norms. My point is not that internalism cannot rule out this kind of justification-undermining luck, but rather that existing internalist views cannot accomplish this task unless they require compliance with the relevant internalist norms *in the modal space*.

This can best be seen by putting the problem that this kind of j-luck poses in terms of the *basing relation*. Internalists typically endorse the claim that doxastic justification requires that beliefs be properly based on evidence or good justifying reasons.[17] As Bondy and Pritchard (2017) argue, subjects may be lucky to form beliefs for which they have good evidence (i.e., propositionally justified beliefs) in cases where they do not believe on the basis of the good evidence available. However, even if we get subjects to believe on the basis of good evidence, this may not be sufficient for doxastic justification. Turri (2010), for example, offers several cases in which basing beliefs on good evidence does not yield doxastic justification. However, these are cases in which the beliefs in question are clearly *improperly* based on good evidence (e.g., believing that p on the basis of a tea-leaf reading on reasons that propositionally justify p).

In contrast, the cases considered here are cases in which the beliefs in question are properly based according to extant internalist views, but would or could easily have been not so. They are cases in which the basing appears to be done in the right way, but in a lucky way that shows that the basing relation is not adequate enough to yield doxastic justification. The case used for illustration is one in which the subject makes errors of reasoning under the same relevant conditions just before and after the actual case in which the subject reasons correctly. The kind of luck involved here forces us to conclude that the subject either did not follow the correct rule of inference, even if it appeared that he did, or that he followed it without really understanding that it was a correct rule or the range of cases to which it applies. So even if the subject gets things right in the actual situation, and even if he does so on the basis of good evidence, he does not base his belief on evidence in the right kind of way. For proper basing (and doxastic justification more generally) cannot arise arbitrarily.

2.2 Accessibilism

---

[17] I use 'evidence' and 'reasons' interchangeably (as e.g. Turri 2010)

Let us begin the critique of internalism with an analysis of its most paradigmatic incarnation: accessibilism. *Accessibilism* is, roughly speaking, the view that being justified in believing p requires that there are factors that contribute to the justification of one's belief that p, and that one has special access to these factors (e.g., through reflection, introspection, memory, a priori reasoning, and so on) (see Pryor 2001 for an overview of accessibilism). More specific versions of this kind of internalist view can be formulated as a function of certain theoretical choices. These include whether the kind of special access to the relevant justificatory factors requires that one be actually or potentially aware of them, whether one must be aware (or potentially aware) of the mere existence of the relevant justificatory factors or also of their adequacy (i.e., the fact that they contribute to the justification of one's belief), and whether being justified in believing p requires one to be justified in being so justified. Below is a list of accessibilist conditions for justification that can be combined in various ways to yield different internalist views of doxastic justification. S's belief that p is justified only if:

(J5) There are factors F that contribute to the justification of the belief that p.

(J6) S is actually aware of F.

(J7) S is potentially aware of F.

(J8) S is aware of the adequacy of F (i.e., of the fact that they contribute to the justification of the belief that p).

(J9) S is justified in believing that S is so justified.

To see how accessibilist justification can be undermined by luck, let us review Adam's story in light of (J5)-(J9). Recall the short interval version of the case, in which Adam forms his beliefs by proper reasoning, but could very easily form those beliefs by improper reasoning. In the short intervals where he reasons badly (i.e., the 'bad' intervals), Adam's beliefs are obviously unjustified. Interestingly, given BAD, Adam's beliefs are not justified even in the 'good' short intervals in which he follows correct rules of inference. The reason, as we have seen, is that it is arbitrary whether he ends up using the correct rules, a point reinforced by the fact that Adam thinks he is inferring correctly at all times. In other words, it is a matter of luck that Adam's beliefs in the 'good' short intervals come about through good reasoning, which undermines their justification. Given this luck, we can ask what accessibilist conditions (J5)-(J9) hold in such intervals. This is important because if all the conditions that an accessibilist view takes to

be individually necessary and jointly sufficient for justification hold in the case under consideration, then that view will not be able to eliminate this kind of dangerous j-luck.

In the good short intervals Adam finds good evidence for p and for p entails q. Moreover, he has evidence (from the philosophy courses he has taken) that applying modus ponens to this evidence leads him to conclude that q. This means that there are factors that contribute to the justification of his belief in q. Therefore, (J5) holds. Adam is also actually aware of the existence of all this evidence, and if for some reason he is not actually aware of it (e.g., because he is distracted), he is certainly potentially aware of it. Therefore, (J6) and (J7) also hold. Moreover, he is also aware of the fact that his evidence is good and that he has applied modus ponens to arrive at his belief in q given the available evidence. Thus, we can assume that he is also aware of the fact that his evidence for p and his evidence for p entails q, as well as his application of modus ponens, all contribute to the justification of his belief in q. This means that (J8) also obtains. Moreover, the prior awareness could easily lead to a higher-order belief that his belief in q is justified in the way described. This higher-order belief could itself be justified if, for example, Adam consults a logic textbook to make sure that modus ponens is a valid rule of inference and that he has correctly followed all the steps in its application. So, as it stands, (J9) also holds. In view of this, not even a very sophisticated accessibilist conception of justified belief that satisfies all conditions (J5)-(J9) would be able to eliminate the dangerous j-luck of this case.

In response to the foregoing, accessibilists might be willing to reject BAD and endorse the counterintuitive result that in the 'good' short intervals in which Adam reasons correctly, his beliefs are justified after all, while in the 'bad' intervals in which he reasons incorrectly, they are not. However, BAD is difficult to reject only in light of (J5)-(J9).

As we have seen, while the specific length of such intervals may need to be different for different people (e.g., one minute, one second long), if the intervals are sufficiently short, it is hard not to conclude that whether Adam complies with the relevant internalist norms is a completely arbitrary matter, especially given that Adam believes he is reasoning correctly in the good intervals but mistakenly believes that he is reasoning correctly in the bad intervals—in such bad intervals he continues to believe that he is following the correct rules of inference (even though he is not). This in turn means that in the good intervals he is unaware (both actually and potentially) of the *modal fact* that he could easily reason, or could have reasoned, incorrectly. Such a modal fact is beyond Adam's ken.

It seems, then, that an accessibilist *anti-j-luck condition* would have to stipulate that one must not violate the relevant internalist requirements in nearby possible worlds, or that one must be actually or potentially aware of such a modal fact, or even of the fact that one does not violate the relevant internalist requirements in nearby possible worlds while thinking that one satisfies them (which is the case with Adam). Whatever the relevant anti-j-luck condition, it will be one that, unlike (J5)-(J9), takes into account whether one's beliefs are justified in the modal environment and not just in the actual situation.[18]

2.3 Evidentialism

This lesson can be extended to another paradigmatic form of internalism, evidentialism, which encounters the same problem for the same reasons. *Evidentialism* is a view that takes a belief to be justified at a time t depending on whether the evidence one has at t fits or supports the proposition believed. That is, by focusing on actual times, evidentialism does not take into account what is happening in the modal environment. In the kind of case we are considering, however, justification is lost in the actual case precisely because one would easily form or could have easily formed a belief in ways that internalists themselves (evidentialists included) consider unsuitable for forming justified beliefs—recall: "beliefs formed by false or arbitrary reasoning are paradigmatic cases of unjustified belief" (Cohen *ibid*.).

If one hour S forms a belief in p based on supporting evidence, but one hour later or earlier S forms or would have formed a belief in p unsupported by evidence (because S unwittingly engages in motivated reasoning), then S's belief in p cannot be justified: it is entirely arbitrary whether S is in a 'good' or a 'bad' hour (reduce the intervals if necessary to agree with this point). Of course, externalists would offer a simple explanation for why S's belief is not justified: the way S forms her belief in p is completely unreliable. But this explanation is not available to

---

[18] Another possible response on behalf of the accessibilist is this: since Adam does not notice the change from good to bad reasoning, and therefore everything seems to be going well from his point of view, the accessibilist might be willing to ascribe justification even in the intervals when Adam is reasoning badly. If so, the kind of luck at stake does not undermine accessibilist justification. However, this move risks running counter the central internalist assumption presented at the beginning of §2.1. In particular, any internalist view should agree with Cohen (*ibid*.) that "[t]here is a fundamental epistemic difference between the beliefs of *A* [the good reasoner] and the beliefs of *B* [the bad reasoner] (...) [which] is marked precisely by the concept of justified belief." The proposed version of accessibilism would not be able to distinguish between good and bad reasoners in evil demon worlds, i.e., in scenarios where reliability is irrelevant.

internalists (accessibilists, evidentialists), not only because it is an externalist explanation that they would never agree with, but also because a stipulation of the case is that Adam lives in an evil demon world and consequently all his beliefs are unreliable anyway. An important lesson for internalists, then, is that while the justification of our beliefs depends largely on the fulfillment of internalist norms that relate to the actual situation, we must impose requirements on the modal environment to preclude dangerous j-luck. Thus, counterfactual conditions for justification seem to be necessary to develop an adequate (internalist) *anti-j-luck epistemology*.

That is the general argument. But let us look at evidentialism a little more closely to see the point more clearly. According to evidentialism (at least according to one of its recent incarnations by Conee and Feldman 2008), S is propositionally justified in believing p at t if and only if S's evidence E at t on balance supports p, where E supports p at t if p is part of the best explanation for E available to S at t.[19] In their seminal 1985 paper, they also offer an account of doxastic justification (or what they call well-founded belief). We can adapt and merge it with their 2008 view of propositional justification as follows. S's belief that p is justified at t if and only if:

> (J10) S's evidence E at t on balance supports p, where E supports p at t if p is part of the best explanation for E available to S at t, and S believes that p on the basis of E.[20]

According to Feldman and Conee (1985), appealing to the basing relation along the lines of (J10) helps distinguish cases in which someone comes to believe that q by modus ponens from p, and p entails q, from cases in which the basis for believing q is not one's evidence for it but the sound of the sentence expressing q. The former is a case of doxastically justified belief, the latter is not, for even if one has relevant good evidence in the latter case, one does not use it to form the belief in question.

Let us revisit Adam's short-interval case in light of (J10). In the 'good' short intervals, Adam comes across evidence for p and for p entails q, and deduces q by modus ponens. On a simple interpretation of the case, Adam's total evidence for q consists in his evidence for p, for p entails q,

---

[19] There are of course other ways to understand the epistemic support relation (see Conee and Feldman 2008 for a useful overview). I use Conee and Feldman's recent view for illustration.

[20] In their 1985 definition, Feldman and Conee include further specifications of what it takes to base a belief on one's evidence to accommodate "the fact that a well-founded attitude need not be based on a person's whole body of evidence" (Feldman and Conee 1985: 33), but these are not relevant to the case at hand. I omit them for simplicity.

and the evidence introduced by the inference itself.[21] The proposition believed, q, is part of the best explanation of Adam's evidence available to him at t. Moreover, Adam's basis for believing q is the earlier good evidence, i.e., it is not as if he believes q because he likes the sound of the sentence expressing q—which is the kind of case they use to motivate the introduction of the basing relation. So there is no reason why we should assume that (J10) does not hold. The problem is that Adam's belief at t meets evidentialist standards but it is *not* justified because of the dangerous epistemic luck in question. For although his evidence at t supports the proposition he believes, and he believes it on the basis of that evidence, Adam could very easily have ignored his evidence and believed q at (t–n) or at (t+n) simply because q agrees with his ideological values, thus arriving at the same belief based on no supporting evidence.

Like accessibilists, evidentialists might try to reject BAD and accept the counterintuitive result that in the 'good' short intervals in which Adam reasons correctly, his beliefs are justified, whereas in the 'bad' intervals in which he reasons incorrectly, they are not. But this is hardly tenable if the alternating time intervals are sufficiently small. In case this is unclear, let us paraphrase Cohen's previous quote and imagine two inhabitants of an evil demon world, A, who is a good reasoner, i.e., who reasons according to the correct rules of inference, and B, who is a random reasoner, i.e., someone who sometimes follows the correct rules of inference and sometimes engages in motivated reasoning, unaware that he is reasoning in this inadequate way. The point, then, is that there is certainly a fundamental epistemic difference between the beliefs of A and B, even in those moments when B happens to form beliefs by good reasoning. This difference has to do with the arbitrariness of B's reasoning, and thus with the luck involved, which leads to a difference in justification. The problem is that, as in the case of accessibilism, evidentialism regards only *actual* factors as relevant to the justification of beliefs, whereas cases of dangerous j-luck show that *modal* factors influence whether an occurrent belief is justified. Thus, although we do well to follow our evidence, this may not be sufficient for justified belief. An open question for evidentialism, then, is how our beliefs in nearby possible worlds should be supported by or based on the evidence, and hence how to avoid licensing beliefs as justified in cases of justification-undermining luck.

---

[21] In this regard, Conee and Feldman (2008: 13) claim that "[p]roperly inferring a proposition from others that are justified is evidence that the inferred proposition is true".

3 J-luck and externalism

One might think that externalists are better positioned to meet the challenge posed by dangerous j-luck. However, this is not obvious. Consider a canonical externalist view of epistemic justification, *process reliabilism*, which holds that S's belief that p is justified if and only if:

> (J11) S's belief that p has been formed by a reliable belief-forming process.[22]

The challenge to process reliabilism is twofold. First, it is unable to rule out the kind of dangerous j-luck involved in cases of fleeting processes, as we shall see in a moment. Second, regardless of what one takes to be a reliable belief-forming process, dangerous j-luck can still occur in evil demon worlds, where beliefs are arguably unreliably formed. However, as I will argue, an adequate theory of doxastic justification should be able to tell the difference between the good and the random reasoner in these worlds, even if one believes that none of their beliefs are justified. Process reliabilism seems incapable of doing this unless it is further qualified.

One such qualification comes from externalist forms of virtue epistemology, such as *virtue reliabilism*, according to which justified beliefs are competently formed beliefs. Virtue reliabilism helps with the first problem, but not so clearly with the second.

The difference between an epistemic competence (or a cognitive virtue or ability) and a merely reliable belief-forming process is that the former is (1) integrated in one's cognitive character and that (2) it exhibits a kind of stability that the latter need not. (1) and (2) are best illustrated by the problem of strange and fleeting processes. Here is John Greco (2010):

> Reliabilist theories have long been plagued by counterexamples involving strange and fleeting cognitive processes. The idea is that if a cognitive process is either strange enough or fleeting enough, then it will not give rise to knowledge even if it is reliable. (Greco 2010: 149)

Greco gives the following well-known example of a *strange cognitive process*:

> Consider the case of the Serendipitous Brain lesion. Suppose that S has a rare brain lesion, one effect of which is to reliably cause the true belief that one has a brain lesion. Even if the process is perfectly reliable, it seems wrong that one can come to have

---

[22] See Goldman (1979) for the original formulation of this view.

knowledge that one has a brain lesion on this basis. (Greco 2010: 149)

And the following example of a *fleeting process:*

> Consider the case of the careless Math Student. Suppose that S is taking a math test and adopts a correct algorithm for solving a problem. But suppose that S has no understanding that the algorithm is the correct one to use for this problem. Rather, S chooses it on a whim, but could just as well have chosen one that is incorrect. By hypothesis, the algorithm is the right one, and so using it to solve the problem constitutes a reliable process. It seems wrong to say that S thereby knows the answer to the problem, however. (Greco 2010: 149)[23]

What is true of knowledge is also true of justification: in neither case are the beliefs in question justified. The case of interest for our purposes is that of the math student, whose method of belief formation is reliable but fleeting. The reason is that it is a case of *dangerous j-luck* (insofar as epistemic justification is understood in terms of process reliabilism). To see this, just consider that it is pure luck that S uses a reliable method (the right algorithm) and not an unreliable one (a wrong algorithm). In this way, dangerous j-luck (i.e., dangerous to a process-reliabilist conception of justification) arises when it is a matter of luck that the belief-forming process one uses is reliable. The case of the math student, and cases of fleeting processes in general, are thus the externalist counterpart of Adam's short-interval case for internalism.

Of course, internalists have no problem with the math student case, since they can argue that because the student does not understand that the algorithm he is using is the correct one, he does not meet the relevant internalist standards such as access to justificatory factors or support by evidence. However, this explanation is not available to process reliabilists, who are only interested in the reliability of the relevant belief-forming process. But as the case shows, it may be a matter of luck that one's belief-forming process turns out to be reliable.

A possible response on behalf of the process reliabilist would be to consider the choice of the method itself as part of the relevant method of belief formation. This would help address cases of fleeting processes. The problem with this approach is that it risks mixing two different notions of

---

[23] This is different from the case we saw in §2.1 of the math student who answers correctly one out of 200 in an exam. In the latter case, the student arrives at the correct answer by a reliable and *not fleeting* process, which explains why her answer is justified.

reliability: namely, reliability of belief formation in the unconditional standard sense (i.e., high ratio of true to total beliefs) and reliability of method selection (i.e., high ratio of selection of a reliable method of belief formation to total attempts to select a method of belief formation). However, the two can diverge: an agent can exhibit low reliability in method selection and high reliability in belief formation (and vice versa). If justification for belief partially depends on reliability in method choice, that agent's beliefs would not be justified, which would be a wrong result.

Virtue reliabilists are in a better position to address the problem of fleeting processes. The solution they offer (in the case of knowledge), as we have seen, is to appeal to cognitive virtues rather than merely reliable cognitive processes. Not surprisingly, this kind of solution also paves the way for them to address the problem of dangerous j-luck. After all, cases of fleeting process are not just cases of ignorance, but also cases of unjustified belief due to the presence of dangerous epistemic luck, and what is true for knowledge is also true for justification. To illustrate, here is Greco's solution for the case of knowledge:

> [C]ognitive virtues [cannot] be fleeting, for virtues are by nature stable. More exactly, virtues must be stable across close possible worlds, and so get just the sort of modal analysis that is needed to address counterexamples involving fleeting processes. Consider once again the case of the careless Math Student. Plausibly, the student does not have knowledge precisely because there are nearby worlds where he does not use the correct algorithm. The problem is not that the algorithm is unreliable – it is perfectly reliable. The problem, rather, is that the *student* is not reliable – his use of the algorithm is not grounded in a cognitive ability, and so there are close possible worlds where he chooses the wrong algorithm instead of the right one. (Greco 2010: 150)

In his solution, Greco assumes a modalized notion of cognitive virtue. Although the formulation of such a notion is controversial, Greco suggests one way in which externalists might address the problem of eliminating dangerous j-luck, namely by introducing *modal conditions* for justification. Here is one of them for illustration. S's belief that p is justified only if:

> (J12) S's belief that p is the product of an exercise of S's cognitive virtues in conditions C, where S has a cognitive ability relative to C

only if across close possible worlds where S is in C, S has a high rate of success in forming true rather than false beliefs.[24]

The *externalist anti-j-luck condition* for justification need not be this one. Virtue reliabilists might prefer other modalized conceptions of cognitive virtue, and hence other modal conditions for justification. The general point remains: as in the case of internalism, modal conditions for justification seem to be necessary to preclude dangerous j-luck.

However, the prospects for externalist anti-j-luck modal conditions are dimmer than one might think. This is because such conditions provide only a partial solution to the problem of eliminating justification-undermining luck. In particular, they only serve to exclude cases of fleeting processes as cases of justified belief. Externalists, however, should also account for, or at least explain away, the difference in justification between the good reasoner and the random reasoner in evil demon worlds. In such worlds, neither the good nor the random reasoner is cognitively virtuous, because their success rate in forming true rather than false beliefs is equally bad in both the actual and nearby possible worlds. After all, the demon is evil enough to ensure that *all* their beliefs become false. So their beliefs are not justified according to (J12) or similar conditions. But there is a difference in justification between good and random reasoners in evil demon worlds that must either be explained or at least explained away by some distinction. To deny that there is nothing to explain by simply appealing to the fact that all their beliefs are unreliable and therefore not justified is just as unsatisfactory as the outright denial that the new evil demon problem is a real problem. But *it is* a problem that many externalists have taken seriously. Thus, to explain why the beliefs of random reasoners are not justified due to luck, externalists may need to resort not to modal conditions for justification, but to the same sort of maneuvers they use in answering the new evil demon problem. For example, they can distinguish a sense of justification according to which a belief is justified only if it is produced by a cognitive virtue that would be reliable relative to non-demon worlds but unreliable in demon worlds (e.g., Sosa 2003). Or they may distinguish the externalistically justified beliefs of us, the inhabitants of the real world, from the blameless (or excused) unjustified beliefs of good reasoners in evil demon worlds, and these in turn from the blameworthy unjustified beliefs of bad reasoners in those worlds (e.g., Williamson *forthcoming*).

---

[24] This condition is inspired by Greco's account of ability possession (Greco 2010: 77).

The bottom line is that the problem of eliminating dangerous j-luck, though not an intractable problem, becomes somewhat more difficult for externalists than for internalists—just as eliminating dangerous k-luck is more difficult for internalists than for externalists. None of them is off the hook, however: justification-undermining luck remains a problem for everyone.

4 Concluding remarks

A widespread assumption in epistemology is that epistemic luck is primarily a problem for theories of knowledge. In this paper, however, I have argued that epistemic luck is a more pervasive phenomenon than is commonly assumed. It occurs not only in the theory of knowledge but also in the theory of justification and, in particular, of doxastic justification. As in the case of knowledge, beliefs, I have argued, can fall short of justification due to luck. Moreover, the problem of eliminating dangerous j-luck is a problem for both externalist and internalist views of justification, although it may be more difficult for externalists. Certainly there are ways to solve the problem, and we have seen some suggestions for how externalists and internalists might modify their views to eliminate justification-undermining luck. But the project of this paper was not to solve the problem, but to provide an explanation of what the problem is. That, I believe, has been accomplished.

References

Aspernäs, J., Erlandsson, A., & Nilsson, A. (2022). Motivated Formal Reasoning: Ideological Belief Bias in Syllogistic Reasoning across Diverse Political Issues. *Thinking & Reasoning*. https://doi.org/10.1080/13546783.2022.2038268

Baumann, P. (2012). No Luck with Knowledge? On a Dogma of Epistemology. *Philosophy and Phenomenological Research* 89: 523-51.

Bondy, P. & Pritchard, D. (2017). Propositional Epistemic Luck, Epistemic Risk, and Epistemic Justification. *Synthese* 195: 3811-20.

Calvillo, D. P., Swan, A. B., & Rutchick, A. M. (2020). Ideological Belief Bias with Political Syllogisms. *Thinking & Reasoning* 26: 291–310.

Chisholm, R. M. (1977). *Theory of Knowledge*. Prentice-Hall.

Coffman, E. J. (2009). Does Luck Exclude Control? *Australasian Journal of Philosophy* 87: 499-504.

Cohen, S. (1984). Justification and truth. *Philosophical Studies* 46: 279-95.

Conee, E., & Feldman, R. (2008). Evidence. In Q. Smith (ed.), *Epistemology: New Essays*. Oxford University Press.

de Grefte, J. (2018). Epistemic Justification and Epistemic Luck. *Synthese* 195: 3821-36.

Engel, M. (1992). Is Epistemic Luck Compatible With Knowledge? *Southern Journal of Philosophy* 30: 59-75.

Evans, J. S. B., & Curtis-Holmes, J. (2005). Rapid Responding Increases Belief Bias: Evidence for the Dual-Process Theory of Reasoning. *Thinking & Reasoning* 11: 382–389

Feldman, R. & Conee, E. (1985). Evidentialism. *Philosophical Studies* 48: 15-34.

Gampa, A., Wojcik, S. P., Motyl, M., Nosek, B. A., & Ditto, P. H. (2019). (Ideo)logical Reasoning: Ideology Impairs Sound Reasoning. *Social Psychological and Personality Science* 10: 1075–1083.

Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis* 23: 121-3.

Goldman, A. (1976) Discrimination and Perceptual Knowledge. Journal of Philosophy 73: 771-791.

—— (1979). What is Justified Belief. In G. Pappas (ed.), *Justification and Knowledg*e. D. Reidel.

Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity.* Cambridge University Press.

Hazlett, A. (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research* 80: 497-522.

Lando, T. (2017). Coincidence and Common Cause. *Noûs* 51: 132-51.

Madison, BJC (2010). Epistemic Internalism. *Philosophy Compass* 5: 840-53.

Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.

Owens, D. (1992). *Causes and Coincidences*. Cambridge University Press.

Pritchard, D. (2005). *Epistemic Luck*. Oxford University Press.

——(2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109: 247-79.

—— (2014). The Modal Account of Luck. *Metaphilosophy* 45: 594-619.

Pryor, J. (2001). Highlights of recent epistemology. *British Journal for the Philosophy of Science* 52: 95-124.

Sosa, E. 2003. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues.* Blackwell.

Sutton, J. (2007). Stick To What You Know. *Noûs* 39: 359-96.

Turri, J.; Buckwalter, W. & Blouw, P. (2015). Knowledge and Luck. *Psychonomic Bulletin and Review* 22: 378-390.

Unger, P. (1968). An Analysis of Factual Knowledge. *Journal of Philosophy* 65: 157-70.

Vahid, H. (2001). Knowledge and Varieties of Epistemic Luck. *Dialectica* 55: 351-62.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.

——(*forthcoming*) Justifications, Excuses, and Sceptical Scenarios. In F. Dorsch & J. Dutant (eds.), *The New Evil Demon: New Essays on Knowledge, Rationality and Justification.* Oxford University Press.