Evidence, Epistemic Luck, Reliability, and Knowledge

Mylan Engel Jr.[1]

Abstract
In this article, I develop and defend a version of reliabilism – *internal reasons reliabilism* – that resolves the paradox of epistemic luck, solves the Gettier problem by ruling out veritic luck, is immune to the generality problem, resolves the internalism/externalism controversy, and preserves epistemic closure.

Keywords Epistemic luck · Reliabilism · Generality problem · Internalism/externalism debate · Personal and doxastic justification · Analysis of knowledge

1 A Modest Goal

My goal is to develop and defend a version of reliabilism—*internal reasons reliabilism*—that resolves the paradox of epistemic luck, solves the Gettier problem by ruling out veritic luck, is immune to the generality problem, resolves the internalism/externalism controversy, and preserves epistemic closure. Let's begin!

2 The Epistemic Luck Paradox

Epistemic luck is a generic notion used to describe various ways in which it is somehow *accidental*, *coincidental*, or *fortuitous* that a person has a true belief that *p*. The phenomenon of epistemic luck gives rise to an epistemological paradox. The paradox is generated by three extremely plausible theses.

2.1 The Knowledge Thesis

We know a lot. We possess all sorts of knowledge about the world around us. You know that you are currently reading an article on epistemology. I know that I am looking at a computer screen. You know what city you are currently in. I know that

---

✉ Mylan Engel Jr.
mylan-engel@niu.edu

[1] Department of Philosophy, Northern Illinois University, DeKalb, IL 60115, USA

I'm drinking a cup of coffee. We both know that whales are mammals and that water is H_2O. We also know that the Trump administration's handling of the COVID-19 pandemic was woefully inadequate. We know these things and so much more.

2.2 The Incompatibility Thesis

A tradition dating back to Plato has taken epistemic luck *simpliciter* to be incompatible with knowledge. Epistemologists have long held that knowledge requires more than mere true belief because a person can fortuitously arrive at a true belief by guessing and lucky guesses fall short of knowledge. Indeed, the role of the justification condition in the justified-true-belief (JTB) analysis of knowledge is to rule out lucky guesses as instances of knowledge. Even so, on fallibilistic accounts of justification, it is possible to be justified in holding a belief that only luckily turns out to be true, which Edmund Gettier (1963) demonstrated with the following example: Smith justifiably believes that Jones owns a Ford [*F*] based on good, but misleading, evidence. On the basis of her justified-but-false belief that *F*, Smith justifiably infers that *F or B* (where *B* = Brown is in Barcelona), despite having no idea of Brown's whereabouts. Coincidentally, Brown is in Barcelona. So, it's just a matter of luck that Smith's justified belief that *F or B* is true. Since Smith's *luck-infused* justified-true-belief that *F or B* clearly falls short of knowledge, fallibilistic justification is incapable of eliminating all forms of knowledge-destroying luck.

2.3 The Ubiquity Thesis

Epistemic luck is an all-pervasive phenomenon that infects every epistemology in one form or other. Its inescapability can be demonstrated as follows: to convert true belief to knowledge, every viable epistemology requires satisfying either some internalistic justification condition or some externalistic condition (that may or may not be a justification condition). But neither an internalistic nor an externalistic condition can completely succeed in eliminating epistemic luck. A little recognized consequence of the new devil demon problem is that internalistic justification is not conceptually connected to truth in any robust way, for demon-world victims have internalistically justified beliefs almost all of which are false. Given the absence of a robust truth connection, it's always in some sense a matter of luck when a *merely* internalistically justified belief turns out to be true. Truth-connected externalist approaches (e.g., reliabilist, truth-tracking, and safety-based accounts) avoid this kind of epistemic luck. However, they are subject to another kind of ineliminable epistemic luck. Consider, e.g., the externalist condition of being a reliably produced belief. While a reliably produced belief's being true is not a matter of luck, one's having reliable cognitive faculties is a matter of luck, for in a demon world, those same cognitive faculties are highly unreliable. Or consider the externalist condition of being a *safe* belief (to be explained below). While a safe belief's being true isn't lucky, having safe beliefs is, for in a demon world, one's beliefs are not safe. Analogous considerations can be applied to any externalist condition of knowledge. Since every epistemology incorporates either an internalistic justification condition or an

externalistic condition, no epistemology can rid us of epistemic luck's intractable presence.

2.4 The Skeptical Challenge

Epistemic luck, then, is ubiquitous and unavoidable. If all forms of epistemic luck are incompatible with knowledge, as the incompatibility thesis maintains, then skepticism is correct, and the knowledge thesis is false. And yet, we remain convinced that we possess lots of knowledge. The task facing the anti-skeptical epistemologist is to reconcile the rather strong intuition that epistemic luck is not compatible with knowledge with the equally evident assumption that it must be. Since the ubiquity thesis is unassailable, the anti-skeptical epistemologist must reject the incompatibility thesis.

3 Rejecting the Incompatibility Thesis

Peter Unger (1968) was the first epistemologist to note that not all forms of epistemic luck are incompatible with knowledge. He identified the following three types of benign epistemic luck:

(1) *Propositional luck*: It can be entirely accidental that *p*, and *S* can still know that *p*. For example, a person who witnesses an automobile accident can certainly know that the accident occurred.

(2) *Existential luck*: For *S* to know that *p*, *S* must exist, and it might be extraordinarily lucky that *S* exists. If *S* is the lone survivor of a fiery plane crash, *S* is lucky to be alive, but *S*'s existential luck does not preclude her from knowing that she survived the crash.

(3) *Facultative luck*: To know that *p*, *S* must possess the cognitive skills requisite for knowledge. Suppose *S* is shot in the head and the bullet narrowly misses all vital regions of the brain required for conceptual thought and knowledge. *S* is overwhelmingly lucky that she still possesses the cognitive capacities needed for knowledge, but since she does possess them, she is still capable of knowing many things, including that she was shot in the head. Unger has successfully identified three types of harmless epistemic luck, but not all forms of epistemic luck are benign. What we need is an account of knowledge-undermining luck.

4 Knowledge-Destroying Epistemic Luck

In "Is Epistemic Luck Compatible with Knowledge?" (Engel, 1992a) (hereafter IELCK), I distinguished two kinds of epistemic luck—*evidential luck* and *veritic luck*—and argued that only veritic luck is incompatible with knowledge. There, I characterized these two types of luck as follows:

EL   A person *S* is *evidentially lucky* in believing that *p* in circumstances *C* iff it is just a matter of luck that *S* has the evidence *e* for *p* that she does have, but given *e*, it is not a matter of luck that her belief that *p* is true in *C*. (Engel, 1992a, 67)

VL   A person *S* is *veritically lucky* in believing that *p* in circumstances *C* iff, given *S*'s evidence for *p*, it is just a matter of luck that *S*'s belief that *p* is true in *C*. (Engel, 1992a, 67)

To see that evidential luck is compatible with knowledge, suppose that a bank robber's mask slips momentarily during a holdup and the startled teller sees clearly that the robber is the bank president.[1] In such a situation, the teller would clearly be lucky to have the evidence she does, but she would nevertheless know that the bank president is the villain.

I contend that all genuine Gettier cases involve veritic luck. In Gettier's original example presented above, Smith's belief that *F or B* is veritically lucky: given Smith's misleading evidence of Jones's Ford ownership status and her total lack of evidence concerning Brown's whereabouts, it's just a matter of luck that Smith's belief that *F or B* is true. Veritic luck with respect to *p* is incompatible with knowing that *p*, because it undercuts the connection between *S*'s evidence for *p* and *p* in a way that makes it entirely coincidental from *S*'s perspective that *p* is true. It's one thing to identify veritic luck as the kind of epistemic luck that destroys knowledge. It's another to spell out the condition that rules out veritic luck. In IELCK, I only did the former. In Section 9, I will do the latter.

Duncan Pritchard (2003) agrees that, of these two types of luck, only veritic luck is incompatible with knowledge, but he replaces my evidence-based characterization of veritic luck with the following modal analysis:

MVL   For all agents *S* and propositions *p*, the truth of *S*'s belief that *p* is veritically lucky iff *S*'s belief that *p* is true in the actual world α but false in nearly all nearby possible worlds in which S forms the belief in the same manner as in α. (Pritchard, 2003, 111)

MVL differs from VL in the following way: it concerns the connection between the *method of belief formation* and proposition believed, rather than the connection between *S*'s *evidence* and the proposition for which it is evidence. Pritchard argues that a safety-based neo-Moorean account, according to which knowledge is *safe true belief*, is capable of eliminating veritic luck. In a moment, we will see, *contra* Pritchard, that safe true belief is incapable of ruling out certain paradigm cases of veritic luck.

4.1 Attempts to Block Veritic Luck

The post-Gettier literature is rife with attempts at supplementing or amending the traditional JTB analysis with a satisfactory anti-luck constraint on knowledge. The

---

[1] For the original discussion of the bank-robber-mask-slipping case, see Nozick 1981.

first wave of proposals includes adding a no-false-grounds constraint to JTB (Myers & Stern, 1973), replacing JTB with a causal theory of knowing (Goldman, 1967), and supplementing JTB with a no-defeaters condition (Chisholm, 1964; Lehrer & Paxson, 1969; and Klein, 1971). These and similar proposals fall prey to ever-more-complicated Gettier cases. The general consensus is that none of these proposals succeeds. Second-wave luck-eliminating proposals invoke counterfactual or sub-junctive constraints on knowing, principal among them: *sensitivity* and *safety*.

4.1.1 Sensitivity

Robert Nozick (1981) argues that knowledge requires *sensitive belief*, where sensitive belief is defined as follows: *S*'s belief that *p* is *sensitive* iff were *p* false, *S* would not believe that *p* (i.e., iff *S* does not believe *p* in the closest ~*p*-world). To be sure, sensitive belief does preclude veritic luck, but it does so at a steep price. First, Nozick's sensitive-true-belief account of knowledge results in closure failure.[2] On Nozick's account, I know that I have hands because my true belief that I have hands is sensitive. (If I didn't have hands, I wouldn't believe that I have hands.) I also know that my having hands entails that I am not a handless brain in a vat (BIV) deceived into thinking I have hands. But, on the sensitive-true-belief account, I don't know that I'm not a handless BIV because my belief that I'm not a handless BIV is not sensitive. (If I were a handless BIV, I would still believe I'm not a handless BIV). In short, Nozick's sensitive-true-belief account of knowledge entails the following abominable conjunction: I know that I have *hands*, but I don't know that I'm not a *handless* BIV.[3]

Second and even more problematic, knowledge does not require sensitivity, as Jonathan Vogel (1999) has clearly demonstrated with *hole in one*, a variation of which runs as follows: the fourth hole at Augusta National Golf Course where The Masters is played is a tricky 240-yd par 3, euphemistically called "Flowering Crab Apple." In the 2021 Masters, not one player shot a hole in one on this diabolical hole, and there were only 9 birdies throughout four rounds of play.[4] Right now, I *know* that not all seventy-two players in the 2022 Masters will shoot a hole in one on Flowering Crab Apple in the first round of play, but my belief to this effect is not sensitive. If every golfer were to shoot a hole in one on Flowering Crab Apple in Round One of the 2022 Masters in defiance of the astronomical odds against such an event, I would still believe that they weren't going to do so. So, sensitivity is not necessary for knowledge.

---

[2] According to the principle of epistemic closure, if I know that *p* and I know that *p* entails *q*, then I know (or at least am in a position to know) that *q*.

[3] I follow Keith DeRose in referring to such egregious instances of closure failure as "abominable conjunctions" (1995, 27–29).

[4] To put the difficulty of Flowering Crab Apple into perspective, consider this: in the 2021 Masters, there were 9 times more bogeys than birdies—81 bogeys in all—throughout four rounds of play.

4.1.2 Safety

Considerations such as these have led a number of epistemologists (Sosa, 1999 & 2000, Williamson, 2000a & 2000b, and Pritchard, 2005) to replace the sensitivity condition with some sort of safety condition. Safety comes in different strengths: *S*'s belief that *p* is *strongly safe* iff were *S* to believe that *p*, *p* would be true, i.e., in all the closest *Bp* worlds, *p* is true. *S*'s belief that *p* is *weakly safe* iff *S* would not easily be mistaken with respect to *p*; i.e., *p* is true in the overwhelming majority of close *Bp* worlds.

Peter Murphy (2005) uses Kripke's famous Nozick-bashing counterexample to show that strong safety results in closure failure. Suppose the following is true of Kripke County: the landscape is peppered with barn façades, there are a few real barns in the county, some of the real barns are red, and some are blue, but all the façades are red. Driving through Kripke County, Henrietta is unaware that most of the barn-looking structures are façades. She looks at a blue barn and comes to believe that she is looking at a blue barn. Her belief is safe. In all nearby worlds, where she believes she is looking at a *blue barn*, she is looking at a blue barn. (There are no blue façades!). However, her belief that she is looking at a *barn* is not safe. There are many nearby worlds where she believes she's looking at a barn but is really just looking at a façade. So, strong safety entails that Henrietta knows she's looking at a *blue barn*, but doesn't know she's looking at a *barn*. Another abominable conjunction.

Avram Hiller and Ram Neta (2007) convincingly argue that no *safe belief* condition can eliminate all cases of veritic luck as follows: start with a justified but false and unsafe belief like Smith's belief that Jones owns a Ford [*F*]. Next, have Smith justifiably infer a random disjunction of the form *F or G*, where unbeknownst to Smith, *G* is true in all nearby worlds. For example, let *G* = Guinevere will not win a Grammy. Suppose that, unbeknownst to Smith, Guinevere is totally devoid of musical talent and there is no remotely close world where she wins a Grammy. Then, Smith's justified true belief that *F or G* will be safe but veritically lucky nonetheless.

4.2 Diagnosis

The reason safe true belief accounts are incapable of ruling out veritic luck is because they fail to appreciate the fact that knowledge-destroying luck affects the connection between one's evidence for *p* and *p*, not the modal relation between one's belief that *p* and *p*. Thus, MVL is misguided. It is VL that correctly characterizes the kind of veritic luck that destroys knowledge. In order to solve the Gettier problem, we must find a condition that rules out veritic luck so understood. In Sec. 9, I will provide such a condition.

5 Reliabilism and the Generality Problem

Despite its popularity among externalists, process reliabilism faces a potentially devastating problem known as *the generality problem*. To explain why this problem arises, I need to say more about the notion of a belief-forming cognitive

process. According to process reliabilism, as per Alvin Goldman (1979), a belief-forming cognitive process is a *functional operation or procedure* that generates a *mapping* from certain input states to other output states. The inputs are perceptual experiences, belief states, introspective states, seemings, and memorial states. The outputs are belief states, i.e., states of believing this or that proposition. A *token* process is a single instance of such an input–output operation. A process *type* can be repeatedly instantiated by different process tokens. Process reliabilism identifies processes with *types* rather than *tokens*, since only types have statistical properties like producing truth 90% of the time, and it is these sorts of statistical properties that determine the reliability of a process. Of course, types are abstract entities and are incapable of being causes, and yet, process reliabilists maintain that processes *cause* beliefs. To avoid the causal inefficacy problem, process reliabilism maintains that it is process tokens (token instances of a process type) that *cause* beliefs, but it is the reliability of the process type (instantiated by a given process token) that determines the degree of justification of the resulting belief (Goldman, 1979, 11).

We are now in a position to see why the generality problem arises for process reliabilism. Process types can be described with a greater or lesser degree of *generality*, and the degree of generality will partly determine the reliability of the process. Any given process token will be an instance of indefinitely many process types—types with differing degrees of generality and reliability. For example, the token perceptual process that produced my current belief that there is an indigo bunting at the birdfeeder is an instance of each of the following process types: my vision, human vision, mammalian vision, my vision at a particular time of day, my vision in a particular year, my vision in a particular century, my vision with glasses, and my vision at a distance. Each of these process types has a different degree of reliability. Suppose that human vision in general is reliable but that my vision happens to be extremely poor and unreliable. Also suppose that my vision with glasses is reliable but my vision of distant objects is unreliable. Is my belief justified because it is produced by human vision, or is it unjustified because it is produced by my vision? It is justified because it is produced by my vision with glasses or unjustified because it is produced by my vision at a distance? Without an account of process relevance that specifies the relevant process type, process reliabilism is an empty theory, for without such an account, process reliabilism has no implications about the justification of beliefs. To date, no reliabilist has been able to provide such an account. If the generality problem cannot be solved, *process reliabilism* will have to be abandoned.

There are basically two ways of resolving the generality problem for reliabilism. One could provide criteria for identifying the relevant process type, or one could develop an alternative version of reliabilism that is immune to the problem. I think Earl Conee and Richard Feldman are right: the prospects of doing the former are "hopeless" (1998, 24). In Sec. 9, I will take the second tack and develop a version of reliabilism—*internal reasons reliabilism*—that is immune to the generality problem. Internal reasons reliabilism will allow us to circumvent the generality problem while also unifying a number of seemingly discordant but nevertheless widely embraced epistemological theses.

6 Seven Widely Held Epistemological Theses

Each of the seven theses identified below has been widely endorsed and defended by epistemologists. Despite their widespread endorsement, many readers will feel pressure to reject one or more of these theses due to prior theoretical commitments. Indeed, internalists will typically reject VI and externalists will typically reject III. Nevertheless, I submit that each of these seven epistemological theses is true and that it is a virtue of an epistemological theory to be able to accommodate and embrace all of them.

6.1 Seven Epistemological Theses

I. Epistemic justification is necessary for knowledge (nearly every epistemologist).

II. The kind of epistemic justification needed to convert true belief to knowledge must be conceptually connected to truth (BonJour, 1978; Lehrer & Cohen, 1983).

III. Epistemic justification is a function of internal epistemic reasons, where an *internal epistemic reason* for *p* is a state internal to the believer that evidentially supports the truth of *p* (Pollock, 1986).

IV. There is no conceptual connection between internal justification and truth. So, no internalistic theory of justification can provide the needed conceptual connection between justification and truth (Chisholm, 1986; Engel, 1992a).

V. An externalistic reliabilist account of justification can provide the needed conceptual connection between justification and truth (BonJour, 1998; Engel, 1992a).

VI. In order for a belief to be justified, it must be reliably produced (Goldman, 1979).

VII. Process reliabilism, the most widely accepted externalist theory of justification, is open to the generality problem (Conee & Feldman, 1998).

Let me say a bit more about each of these theses.

6.1.1 Thesis I

The idea that epistemic justification is necessary for knowledge is nearly universally accepted by epistemologists. The reason justification is needed for knowledge is that one can arrive at a true belief simply as a result of a lucky guess, and lucky guesses fall short of knowledge. The role of the justification condition in the theory of knowledge is to rule out lucky guesses. The motivation behind the justification condition lies in the conviction that if a belief is epistemically justified, then it's not just a matter of luck when that belief turns out to be true.

6.1.2 Thesis II

In order for epistemic justification to fulfill its luck-eliminating role in the theory of knowledge, there must be a conceptual connection between justification and truth. Call such a connection *the truth connection*.[5] In order for there to be a conceptual connection between justification and truth, for any world *W*, the conditions that make a belief justified in *W* must ipso facto make that belief objectively probable in *W*. Bracketing Gettier situations, when justification is conceptually connected to truth in this way, it is not just a matter of luck when justified beliefs turn out to be true.

6.1.3 Thesis III

It is a truism among internalists that justification is a function of internal reasons. In fact, many of the supposed counterexamples to reliabilism involve cases where a person is able to reliably form beliefs in a particular area without having any internal reasons for these beliefs, e.g., Laurence BonJour's Norman-the-clairvoyant case (1980, 62) and Keith Lehrer's TrueTemp case (1990, 163f).[6] From the internalistic subjective point of view, since Norman and TrueTemp lack internal reasons for their respective beliefs, they are both unjustified in their respective beliefs, even though their beliefs are produced by processes that are 100% reliable.

---

[5] I borrow this terminology from Keith Lehrer and Stewart Cohen (1983, 191).

[6] BonJour's Norman-the-clairvoyant case runs as follows:

"Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable" (BonJour 1980, 62).

Here is BonJour's assessment of Norman and his belief about the President's whereabouts:

"Is Norman epistemically justified in believing that the President is in New York City, so that his belief is an instance of knowledge? According to the modified externalist position [for our purposes, *process reliabilism*], we must apparently say that he is. But is this the right result? Are there not still sufficient grounds for a charge of subjective irrationality to prevent Norman's being epistemically justified?... Norman's belief about the President's whereabouts is epistemically irrational and irresponsible, and thereby unjustified, whether or not he believes himself to have clairvoyant power, so long as he has no justification for such a belief" (BonJour 1980, 62–63).

Lehrer's asks us to imagine a person named TrueTemp, who undergoes brain surgery, and that during the surgery an experimental surgeon implants a perfectly reliable temperature-detecting device (a tempucomp) in TrueTemp's brain. The device accurately detects the temperature wherever TrueTemp is and causes TrueTemp to form an accurate belief of temperature there. Forming temperature beliefs via the tempucomp device is a completely reliable belief-forming process. In addition, we are told that TrueTemp has no idea that a tempucomp has been inserted in his brain, and despite being puzzled by having so many temperature beliefs, TrueTemp never checks to see whether any of his temperature beliefs are accurate. He just accepts these beliefs unreflectively. Reflecting on TrueTemp's epistemic situation, Lehrer writes: "Thus, he [TrueTemp] thinks and accepts that the temperature is 104 degrees. It is. Does he know that it is? Surely not. He has no idea whether he or his thoughts about the temperature are reliable" (Lehrer 1990, 164).

6.1.4 Thesis IV

That no internalist theory of justification can provide a truth connection can be demonstrated as follows: every internalistic theory of justification maintains that the conditions that justify S in holding a belief are entirely specifiable in terms of states internal to S, and for any set of entirely internally specifiable justification-conferring conditions, C_I, there will always be a possible world, W_D, where a Cartesian evil demon has seen to it that S possesses the requisite internal states and, hence, satisfies C_I, even though all of S's contingent empirical beliefs are false. Since the conditions C_I that make S internalistically justified in believing b in W_D do not make b objectively probable in W_D, no internalist theory is capable of providing a truth connection. Internal reasons per se are not conceptually connected to truth.[7]

6.1.5 Thesis V

While internalistic accounts of justification can't provide the requisite truth connection, externalistic reliabilism can.[8] Externalistic reliabilism hold that a belief is justified iff it is reliably produced. If a belief is reliably produced, then that belief is objectively likely to be true because, by definition, the objective indefinite probability of a reliably produced belief being a true belief is high. So, if *reliable production* is the condition that makes a belief justified, then justification will be conceptually connected to truth.

6.1.6 Thesis VI

Collectively, theses II, IV, and V support thesis VI. The fact that (i) the kind of justification needed for knowledge must be conceptually connected to truth [thesis II], (ii) no internalistic account of justification can provide the requisite truth connection [thesis IV], and (iii) an externalistic reliabilist account of justification can provide the needed truth connection [thesis V] suggests that in order for a belief to be epistemically justified, it must be reliably produced.

---

[7] Foreshadowing: in the demon world, internalistic reasons are not reliable indicators of the truth, and so, on the view that I will be defending, these reasons would fail to confer justification on the beliefs they produce in the demon world (per thesis II). However, since these same reasons are reliable indicators in the actual world, they do confer knowledge-generating justification on *our* beliefs in. Thus, unlike purely internalistic approaches, the version of reliabilism I will develop will be able to provide the requisite truth connection while acknowledging the necessity (but nonsufficiency) of internalistic reasons for knowledge-generating justification.

[8] Remember, a *truth connection* is a *conceptual* connection between justification and truth, not merely a *contingent* or *accidental* connection between justification and truth. In order for there to be a conceptual connection between justification and truth, the following condition must obtain: for any possible world W, if conditions C make S's belief that p justified in W, then conditions C make p objectively probable in W.

6.1.7 Thesis VII

Despite the intuitiveness of tying epistemic justification to reliable production (as just noted), internalists insist that process reliabilism succumbs to the generality problem in that there is in principle no non-ad hoc way of identifying the relevant process type for a given process token. Internalists have concluded from thesis VII that reliabilism per se (the view that a belief is justified if and only if reliably produced) must therefore be rejected. Such a conclusion is unwarranted because it is open to the reliabilist to develop an alternative version of reliabilism that does not tie justification to the reliability of the producing process type.

As indicated at the outset, my goal is to solve the epistemic luck paradox, the Gettier problem, and the generality problem in a way that resolves the internalism/externalism controversy. I now turn to that controversy.

7 Internalism vs. Externalism in Epistemology

The internalism/externalism distinction in epistemology is most clearly defined for theories of justification.[9] An internalist theory of epistemic justification is any theory that maintains that epistemic justifiedness is *exclusively* a function of states internal to the cognizer. Externalism is the denial of internalism. Thus, an externalist theory is any theory that maintains that epistemic justifiedness is at least partly a function of states or factors external to the cognizer, i.e., states or factors outside the cognizer's ken.

There is no unified agreement among internalists as to which internal states are epistemically relevant, and different internalisms emerge based on the subset of internal states deemed relevant. Internalists typically maintain that justification is a normative notion in the belief-guiding/regulative sense in that it tells us what we should and should not believe. To perform this belief-guiding regulative function, one must have access to the conditions that justify one in holding a belief. Internalists also typically maintain that one can tell whether one is justified in believing *p* simply by reflecting on one's internal evidence for *p*. The central internalist intuition, as highlighted by the new evil demon problem, is this: there can be no difference in epistemic justification without a difference in epistemically relevant internal states.

---

[9] There is less canonical agreement when it comes to applying the internalist/externalist distinction to theories of knowledge. In one sense, every plausible epistemology is an externalist theory because every plausible epistemology requires an externalist truth condition and an externalist Gettier-blocking constraint. However, in another widely used sense, "externalist" theories of knowledge are theories that replace the internalist justification condition with either an externalist justification condition or some other externalist constraint (such as a causal or modal constraint); while "internalist" theories of knowledge hold that internalist justification is necessary for knowledge and also typically hold that no other kind of justification is needed for knowledge, though they do tend to incorporate some sort of externalist constraint to handle the Gettier problem.

Externalism is motivated by the intuition that epistemic justification must be conceptually connected to truth such that the conditions that make a belief justified in *W* also make it objectively probable in *W*. Externalists are also typically motivated by the view that children and animals can form justified beliefs, while failing to satisfy the internalist's intellectualist requirements for justification. The dominant externalist theory of justification is process reliabilism, a simplified version of which holds that a belief is justified if and only if it's produced by a reliable process.

An adequate resolution of the internalism/externalism controversy must not only identify which conception of epistemic justification—internalism or externalism—is correct; it must also explain why internalists and externalists have such widely divergent intuitions concerning the very same cases (e.g., demon cases, responsible vs. irresponsible demon victims). It is implausible that individuals employing the same concept of epistemic justification would have such radically divergent justification-intuitions regarding the same cases. I submit that internalists and externalists unwittingly employ different conceptions of epistemic justification, that this accounts for their competing intuitions about cases, and that distinguishing these two kinds of justification holds the key to resolving the internalism/externalism controversy. I now turn to that distinction.

8 Personal and Doxastic Justification

The justification requirement for knowledge is usually formulated in one of two ways:

JR1 *S* knows that *p* only if *S* is justified in believing that *p*.
JR2 *S* knows that *p* only if *S*'s **belief** that *p* is justified.

JR1 and JR2 are generally taken to be synonymous formulations of the justification requirement. The synonymy of JR1 and JR2 entails what I call *the equivalency thesis*:

ET *S* is justified in believing that *p* iff *S*'s **belief** that *p* is justified.

Most epistemologists tacitly embrace ET, since they jump back and forth between talking about *S* being justified in a belief and *S*'s belief being justified. However, some are more explicit in their commitment to ET. For example, William Alston writes: "I shall speak indifferently of S's belief that *p* being justified and of S's being justified in believing that *p*" (1989, 83).

I contend that the equivalency thesis is false and that embracing it has led philosophers astray in their epistemological theorizing and criticism. In "Personal and Doxastic Justification in Epistemology" (Engel, 1992b) [hereafter PDJE], I argued that ET is false on the grounds that the two sides of the biconditional have different domains of evaluation. The left-hand side of ET is evaluating the *person S* as being epistemically justified, whereas the right-hand side of ET is evaluating the justificatory status of *S*'s *belief*. There, I dubbed these two kinds of epistemic justification *personal justification* and *doxastic justification*, respectively, and argued that

internalists employ the personal conception of epistemic justification, whereas externalists employ the doxastic conception of epistemic justification.[10] When evaluating whether a *person S* is epistemically justified in believing that *p*, we are assessing whether *S* has formed her belief in an epistemically responsible fashion, given what she had to go on. When evaluating whether a *belief* is epistemically justified, we are concerned with whether or not the belief has a high objective probability of being true. The reason internalists and externalists have competing intuitions about the same cases is because they are evaluating different domains—*persons* and *beliefs*, respectively—as being justified or unjustified. The reason internalists and externalists take themselves to be disagreeing is that they tacitly accept ET and fail to realize that they are employing different concepts of epistemic justification. Once the distinction between personal justification and doxastic justification is recognized, it becomes clear that rather than disagreeing with each other, internalists and externalists have, for the most part, simply been arguing past each other.[11]

In PDJE, I didn't provide a detailed account of personal justification, but I did defend a process reliabilist account of doxastic justification. I still think that the personal/doxastic justification distinction holds the key to the resolving the internalism/externalism controversy, but I am now convinced that it is in principle impossible for process reliabilists to solve the generality problem. I am also convinced that internalists are right in maintaining that epistemic justification in *both* its forms (i.e., personal and doxastic) proceeds in terms of internal reasons. BonJour makes the latter point nicely as follows:

> Knowledge requires epistemic justification, and the distinguishing characteristic of this particular species of justification is . . . its essential or internal relationship to the cognitive goal of truth. Cognitive doings are epistemically

---

[10] As I noted in PDJE, most epistemologists embrace ET and conflate personal and doxastic justification and that has led them astray in their epistemological theorizing and criticism. There are, however, two historically noteworthy exceptions. Alvin Goldman (1979) was the first epistemologist to explicitly distinguish justified *beliefs* from justified *persons*. In "What Is Justified Belief?," he defends a historical process-reliabilist account of justified belief and a terminal-phase reliabilist account of justified persons. He goes on to argue that personal justification is not a kind of justification closely related to knowledge since, to know that *p*, the *entire* causal process that produces the belief that *p* must be reliable/conditionally reliable, not just the final phase of that process. Kent Bach (1985) also recognizes the distinction between justified persons and justified beliefs and argues that internalists employ a responsibilist conception of justified *persons*, whereas externalists maintain that being justified is whatever property a true, ungettiered *belief* must possess to qualify as knowledge. Like Goldman, Bach contends that personal justification has no role to play in the theory of knowledge. As he sees it, knowledge is ungettiered reliably produced true belief. As such, Bach concludes that Norman knows that the President is in New York City, despite having no evidence of the President's whereabouts and despite having no awareness of his clairvoyant ability. Unlike Goldman and Bach, I contend that personal justification *is* necessary for knowledge.

For more recent proponents of the personal/doxastic justification distinction and its importance in epistemology, see Littlejohn 2009 and Koppelberg 2012.

[11] Even after recognizing that internalists employ the personal conception of epistemic justification and externalists employ the doxastic conception of epistemic justification, one substantive dispute remains, namely, which kind of epistemic justification—*personal* or *doxastic*—is the kind of justification required for knowledge. I will address this dispute in Sec. 9.

justified, on this conception, only if and to the extent that they are aimed at this goal—which means roughly that one accepts all and only beliefs which one has *good reason* to think are true. To accept a belief in the absence of such a reason . . . is to neglect the pursuit of truth; such acceptance is, one might say, *epistemically irresponsible*. (BonJour, 1978, 5; first emphasis mine)

The next section draws on the insights gleaned from Secs. 2–8 to develop a coherent comprehensive theory of justification and knowledge that I call *internal reasons reliabilism*.

9 Internal Reasons Reliabilism

Internal reasons reliabilism maintains that epistemic justification (in *both* of its forms) is a function of epistemic reasons. An *epistemic reason* for *p* is a reason to think that *p* is true. It is a consideration that counts in favor of *p*. John Pollock defines an epistemic reason as follows: *R* is an *epistemic reason* for a person *S* to believe that *p* iff it is logically possible for *S* to become epistemically justified in believing *p* by believing it on the basis of *R* (1986, 36). In order for it to be logically possible for *S* to become justified in believing that *p* by believing it on the basis of *R*, *R* must evidentially support or entail *p*. There can be epistemic reasons for *p* that *S* does not possess. However, in order for an epistemic reason *R* to justify *S* in believing that *p*, *S* must possess that reason and must base her belief on it. *S possesses* an epistemic reason to belief that *p* just in case *S possesses* a cognitively accessible internal state that counts in favor of *p* (Kung, 2010, 2). There are doxastic and experiential reasons. A *doxastic reason* for *p* is a cognitively accessible *belief* that counts in favor of *p*. An *experiential* reason for *p* is a cognitively accessible *experience* that counts in favor of *p*. Epistemic reasons can be conclusive or nonconclusive. A *conclusive reason* is a reason that entails that for which it is a reason.[12] *Nonconclusive reasons* rationally support, but do not entail, that for which they are reasons. One important characteristic of nonconclusive reasons is that they are defeasible. A reason *R* for *p* is *defeasible* if the evidence it provides in support of *p* can be defeated by certain other evidential considerations.[13]

As noted above, a person is personally justified in believing that *p* iff she has formed her belief that *p* in an epistemically responsible fashion, and as BonJour rightly notes, believing *p* responsibly requires having good (epistemic) reasons for *p*. Thus, internalists are right in holding that personal justification (the normative, belief-guiding, regulative conception of epistemic justification that they employ) is exclusively a function of internal reasons. Accordingly, we can define personal justification as follows:

---

[12] Here I am following Pollock's understanding of conclusive reasons, as opposed to Dretske's conception of conclusive reasons. On Dretske's conception, a reason *R* for *p* can be "conclusive" even if *R* does not logically entail *p*, provided *R* wouldn't be the case if *p* were false (Dretske, 1971).

[13] For a useful discussion of *defeasible reasoning*, see Pollock 1986, Chapter 2.

9.1 Personal Justification

Personal Justification (PJ): *S* is *personally justified* in believing that *p* iff: (1) *S* bases her belief that *p* on an internal reason *R* for *p*, (2) *R* is either (a) an experiential state of *S* or (b) a belief state of *S*, and (3) *R* is ultimately undefeated for *S* (i.e., *S* possesses no undefeated rebutting, undercutting, or negating defeaters for *R*).[14]

PJ yields the intuitively correct result that my twin in an evil demon world is just as personally justified in believing that he his drinking a cup of coffee as I am. Each of us has the same internal reasons—phenomenologically indistinguishable visual, olfactory, and gustatory sensations—for thinking that we are drinking a delicious cup of freshly brewed Columbian Supremo. However, in my twin's case, the demon is deceiving him by giving him hallucinatory experiences, whereas in my case, presumably, my experiences are veridical. In short, basing one's belief that *p* on an ultimately undefeated internal reason *R* for *p* is sufficient for one to be *personally justified* in believing that *p*.

The little-appreciated flipside of the new evil demon challenge to externalism is that the kind of internalistic personal justification that I share with my demon-world twin lacks a robust conceptual connection to truth, and only a truth-connected conception of justification has any hope of precluding knowledge-destroying epistemic luck. BonJour aptly captures this idea, as well. After telling us that "the idea of being epistemically responsible is the core of the concept of epistemic justification," he immediately continues:

A corollary of this conception of epistemic justification is that a satisfactory defense of a particular standard of justification must consist in showing it to be truth-conducive . . . To insist that a certain belief is epistemically justified, while confessing in the same breath that this fact about it provides no good reason to think that it is true, would be to render nugatory the whole concept of epistemic justification. (BonJour, 1978, 5)

Of course, as BonJour (1998, 127f) himself later recognized, no purely internalistic account of justification can provide the needed truth connection. Only an externalist account of justification can provide a robust conceptual connection between justification and truth. This is where doxastic justification enters the scene.

Let's take stock. *Internalists* are right to insist that epistemic justification (in *both* its forms) requires *internal reasons* (e.g., perceptual experiences, beliefs, introspective experiences, etc.), and they are also right that personal justification is exclusively a function of such reasons. However, *externalists* are right to insist that only

---

[14] Rebutting and undercutting defeaters are well understood, thanks in large part to Pollock's helpful discussion in *Contemporary Theories of Knowledge* (1986, 37–39). A negating defeater can be defined as follows: where *R* is a reason to believe that *p*, *D* is a negating defeater of *R* if *D* is a reason to think that *R* is false.

an externalistic account of epistemic justification can provide the needed conceptual connection between justification and truth. The key to solving the generality problem lies in merging these two ideas. Process reliabilism falls prey to the generality problem because it ties reliability to the producing process type, and any given process token will be an instance of indefinitely many types. The generality problem is circumvented if we tie a *belief's* justificatory status to the reliability of the specific *internal reason* that personally justifies *S* in believing that *p* in the following way: *S*'s *belief* that *p* will be doxastically justified in *W* if and only if it is based on an ultimately undefeated internal reason *R* for *p* that is a *reliable indicator* of *p*'s truth in *W*. Intuitively, *R* is a reliable indicator of *p* iff were *S* to have *R*, *p* would be true. As such, we can think of reliable indicator reasons for *p* as *safe reasons* for *p*. To provide a more precise definition of a *reliable-indicator reason*, I need to introduce the notion of an *R*-world. Let an *R*-world be any world in which *S* possesses reason *R*.

> **Reliable-Indicator Reasons Defined:** Reason *R* is *reliable indicator* of *p* iff all or very nearly all close *R* worlds are *p* worlds.

With this account of the reliable indicator reasons in hand, we can define doxastic justification as follows:

9.2 Doxastic Justification[15]

> **Doxastic Justification (DJ):** For any world *W*, *S*'s belief that *p* is *doxastically justified* in *W* iff: (1) *S*'s belief that *p* is based on an ultimately undefeated

---

[15] My ***reliable-indicator-internal-reasons***–based account of doxastic justification has affinities with views defended by William Alston (1989) and Juan Comesaña (2010), but it was arrived at independently of their views. Indeed, my view was wholly inspired by Fred Dretske's "Conclusive Reasons" and personal conversations with Dretske. It was these personal conversations with Dretske that convinced me that knowledge requires internal reasons. In note 17, I highlight how my account of doxastic justification differs from Dretske's conclusive reasons account of knowledge. Here I will note some key differences between my account of doxastic justification and the views of Alston and Comesaña.

Alston contends that the ***ground*** of a belief is made up of the internal states (beliefs and experiences) that were actually taken up in the formation of the belief, but to be an ***adequate ground***, the ground must be truth conducive. Elsewhere he claims that to be ***adequate*** the ground must be a reliable indication of the fact believed, though he never provides an account of the reliable-indicator relation. Unlike Alston, I cash out the reliable-indicator relation in terms of ***safe reasons***, where ***S***'s reason ***R*** is a safe reason for ***p*** iff were ***S*** to have ***R***, ***p*** would be true. A much more important difference is that Alston formulates his view as an account of ***personal justification***:

S is J_{eg} [J_{eg}=evaluative grounds justified] in believing that ***p iff*** S's believing that ***p***, as S did, was a good thing from the epistemic point of view, in that S's belief that ***p*** was based on adequate grounds and S lacked sufficient overriding reasons to the contrary (1989, 105–106; bold emphasis mine).

There are two problems with Alston's proposal. Since the ***adequacy*** of grounds is an external feature of grounds, ***S*** can't be held epistemically responsible for basing beliefs on justified internal grounds that aren't adequate. While the actual adequacy of grounds can play a role in doxastic justification, it has no role to play in personal justification. Alston's account specified above also renders the wrong personal justification evaluation of demon victims. Since demon victims' grounds aren't adequate in demon worlds, it follows from Alston's account that demon victims aren't personally justified in their beliefs, but this is the wrong result. ***They*** are just as personally justified in their beliefs as we are in our beliefs.

internal reason *R* for *p*, [16] (2) *R* is a reliable indicator of *p* in *W* (in the *safe reasons* sense defined above), and (3) *R* is either (a) an experiential state of *S* or (b) a belief that *S knows* to be true in *W*.[17]

Earlier I observed that PJ yields the intuitively correct result that my twin in an evil demon world is just as *personally justified* in believing that he is drinking a cup of coffee as I am. Each of us has the same internal reasons—phenomenologically

---

Footnote 15 (continued)

To his credit, Comesaña does formulate his evidentialist reliabilism as an account of doxastic justification. The account runs as follows: A ***belief*** that p by S is justified if and only if: (1) S has evidence e; (2) the belief that p by S is based on e; and either (a) e doesn't include any belief and the type ***producing a belief that p based on evidence e*** is actually reliable; or (b) e includes other beliefs of S, all of those beliefs are justified and the type ***producing a belief that p based on evidence e*** is conditionally actually reliable" (2010, 584).

When responding to the new evil demon problem, Comesaña conflates personal and doxastic justification and resorts to a two-dimensional semantics regarding "actually" according to which "an actually reliable process" might refer to the ***diagonal proposition***, which says that the belief was produced by a process that is reliable in the world where it is believed; or it might refer to the ***horizontal proposition***, which says that the belief was produced by a process that is reliable in whichever world the proposition is considered. Comesaña then responds to the new evil demon problem as follows: "In order to amount to knowledge, a belief must be diagonally justified, but even <u>**subjects**</u> for whom knowledge is impossible (like victims of evil demons) can be justified in the horizontal sense" (2010, 579; emphasis mine). The idea that we have to resort to a complicated two-dimensional semantics to account for the intuition that my twin in the demon world is just as justified in his perceptual beliefs as I am in mine strikes me as wholly implausible. I am not evaluating my twin's justificatory status from the standpoint of what is reliable in my world. I am evaluating my twin's justificatory status from the standpoint of what information he had to go on when forming his beliefs. The reason ***he*** is just as justified in his beliefs as I am in mine is because we have exactly the same internal reasons for our beliefs.

A second worry with evidentialist reliabilism concerns condition 2b. Conditional reliability applies to belief-dependent cognitive processes: a belief-dependent process is conditionally actually reliable when a sufficient proportion of its output-beliefs are true ***given that its input beliefs are true***. The problem with 2b is this: conditionally actually reliable processes are ***only*** reliable when ***all*** of their input-beliefs are true, but condition 2b doesn't require that all of the input beliefs be true. It requires that all of the input-beliefs be justified. Since process reliabilism is a fallibilist account of justification, it is possible for a belief to justified and false. When ***S*** believes that ***p*** as a result of a conditionally actually reliable belief-dependent process operating on a justified but false belief ***f***, the resulting belief that ***p*** will not be reliably produced, and yet, according to evidentialist reliabilism, ***S***'s belief that ***p*** would still be justified, despite its unreliable production. In short, Comesaña's account inherits a flaw that was inherent in Goldman's original (1979) account of process reliabilism. My account of doxastic justification avoids this shortcoming. According to DJ, in order for ***S***'s belief that ***p*** to be doxastically justified on the basis of a doxastic reason ***R***, ***S*** must ***know*** that ***R*** is true.

[16] Condition (1) of (DJ) implies that S is personally justified in believing that p. So, personal justification is necessary for doxastic justification, but not vice versa.

[17] My account of doxastic justification is inspired by and closely related to Dretske's account of *having conclusive reasons*. The principle differences are these: first, Dretske offers his account of *having conclusive reasons* as an account of knowledge, whereas I am offering an account of doxastic justification, which is only one component of knowledge. Second, Dretske requires having sensitive reasons for knowledge, whereas my account of doxastic justification only requires safe reasons. Dretske's insistence that knowledge requires sensitive reasons is what causes his view to result in closure failure. Requiring safe reasons for doxastic justification has no such untoward result.

indistinguishable visual, olfactory, and gustatory sensations—for thinking that we are drinking a delicious cup of coffee. However, in the demon world, my twin's internal experiential reasons for thinking that he is drinking coffee are not reliable indicators of the truth. So, according to DJ, his *belief* that he is drinking a cup of coffee is not *doxastically justified*, just as externalists maintain. In the demon world, W_D, the sorts of sensory experiences that my twin and I share are not reliable indicators of the truth. These experiences do not make it objectively probable that my twin is drinking coffee *in W_D*. DJ ties doxastic justification to a subset of internal reasons. Only those internal reasons that satisfy the externalist *safe reasons constraint* (i.e., were S to have R, p would be true) confer doxastic justification. Since such reasons are conceptually connected to truth, doxastic justification is conceptually connected to truth. With one exception [discussed in note 17], doxastically justified true beliefs are not veritically lucky.

Having seen the importance of distinguishing between personal justification and doxastic justification, one question remains: which kind of epistemic justification is necessary for knowledge? The answer is they both are. Personal justification is required to guarantee that S has formed her belief in an epistemically responsible way. Doxastic justification is required to guarantee that S's belief is objectively probable and to block veritic luck. Accordingly, we can define knowledge as follows:

> **Knowledge (K):** *S* knows that *p* iff: (1) *p* is true, (2) *S* believes that *p*, (3) *S* is personally justified in believing that *p* (on the basis of internal reason *R*), (4) *S*'s belief that *p* is doxastically justified, and (5) *S* is a reliable interpreter of *R* (in the sense that there are no nearby worlds where *S* mistakenly takes *R* to be a reason to believe a *p*-incompatible proposition).[18]

---

[18] Space constraints prevent me from being able to discuss the need for condition (5) in any detail, but condition (5) is needed to rule out situations where S only luckily believes that p despite believing it on the basis of a reliable-indicator reason R for p. That situation could arise if S can't reliably tell that R is a reason for p. Suppose there are many nearby worlds where S mistakenly thinks that R is a reason for a p-incompatible proposition q but just happens to believe p on the basis of R in the actual world. In that case, her resulting belief will be veritically lucky despite having a reliable-indicator reason for p, because, given her evidential situation of not being able to reliably see that R supports p rather than q, it is just a matter of luck that she formed the true belief that p rather than the false belief that q. In short, *knowledge requires reliable-indicator reasons reliably interpreted*.

Requiring that S be a reliable interpreter of her reliable-indicator reason for p is somewhat akin to Sosa's *adroitness* requirement. An adroit reasoner "manifests skill or competence" (Sosa 2007, 23) in reasoning. Moreover, when an epistemic agent S arrives at a true belief that p by basing her belief that p on her adroitly interpreted reliable-indicator reasons for p, her true belief that p will be what Sosa calls *apt*; i.e., she will have successfully arrived at her true belief that p *through* her adroit reasoning. (Sosa 2007, 23).

Condition (5) also addresses the problem of knowing necessary truths. Let N represent a particular necessary truth and let R be the reason that S believes N. Since N is true in all possible worlds, R will be a *strongly safe reason* for N no matter what R is and no matter how logically unrelated R might be to N. (Since N is true in *all* possible worlds, N is true all the closest worlds where S has R.) However, if there are close possible worlds where S has R and believes an N-incompatible proposition on the basis of R, then S is in not a reliable interpreter of R as a reason for N, and as a result, S lacks knowledge of N.

According to K, knowledge is responsibly-held, reliably produced true belief.

Let's see how internal reasons reliabilism (IRR) fares with respect to the seven epistemological theses identified in Sec. 6.

I. Epistemic justification is necessary for knowledge.

IRR agrees. Indeed, IRR requires both personal epistemic justification and doxastic epistemic justification for knowledge.

II. The kind of epistemic justification needed to convert true belief to knowledge must be conceptually connected to truth.

IRR agrees. In order to convert a true belief that *S* is personally justified in believing into knowledge, *S*'s *belief* must be doxastically justified. Since doxastic justification is a function of *safe reasons* and safe reasons are conceptually connected to truth, doxastic justification provides the truth connection needed to convert true belief to knowledge.

III. Epistemic justification is a function of internal epistemic reasons, where an internal epistemic reason for p is a state internal to the believer that evidentially supports the truth of p.

IRR agrees. IRR maintains that both personal and doxastic justification are a function of internal reasons.

IV. There is no conceptual connection between internal justification and truth. So, no internalistic theory of justification can provide the needed conceptual connection between justification and truth.

IRR agrees. Personal justification, the kind of purely internalistic justification that internalists embrace, is not conceptually connected to truth, as is demonstrated by the new evil demon problem.

V. An externalistic reliabilist account of justification can provide the needed conceptual connection between justification and truth.

IRR agrees. Doxastic justification as defined in DJ does provide a conceptual connection between justification and truth. According to DJ, the condition that makes *S*'s *belief* that *p* doxastically justified in a world *W* is the condition of *S's basing the belief that p on an ultimately undefeated internal reason that is a reliable indicator of p in W*. Beliefs that are based on internal reasons that are reliable indicators of their truth in *W* have a high objective probability of truth in *W*. Since the condition that makes a belief doxastically justified in *W* also makes that belief objectively probable in *W*, doxastic justification is conceptually connected to truth.

VI. In order for a belief to be justified, it must be reliably produced.

IRR agrees. The kind of epistemic justification that attaches to *beliefs* is doxastic justification. According to DJ, a belief is doxastically justified only if it is based on an internal reason that is a reliable indicator of *p*. Since beliefs that are based on reliable-indicator reasons are ipso facto reliably produced, DJ entails that a belief is doxastically justified only if it is reliably produced.

VII. Process reliabilism, the most widely accepted externalist theory of justification, is open to the generality problem.

IRR agrees. While process reliabilism is open to the generality problem, IRR's reliabilist account of doxastic justification is immune to the generality problem because DJ ties doxastic justification to the *specific internal reason R* that personally justifies *S* in believing that *p*. If *that very reason R* is a reliable indicator of *p* in

W, then *S*'s *belief* that *p* is doxastically justified in *W*.[19] If *that very reason R* is not a reliable indicator of *p* in *W*, then *S*'s *belief* that *p* is not doxastically justified in *W* (no matter how well *R* personally justifies *S* in believing that *p*).

10 Assessing Internal Reasons Reliabilism

As we have just seen, internal reasons reliabilism has the virtue of being able to accommodate and embrace all seven of the widely held epistemologically theses identified in Sec. 6. While a thoroughgoing defense of internal reasons reliabilism must await another occasion, the fact that K yields the right results with respect to every case we have discussed certainly counts in favor of K (and of IRR more generally). In Gettier's original example, Smith's reason for believing that either Jones owns a Ford or Brown is in Barcelona (*F or B*) is Smith's justified-but-false belief that Jones owns a Ford [*F*]. Since *F* is false, it is not a reliable indicator of *F or B*. So, while *Smith* is personally justified in believing that *F or B*, Smith's *belief* that *F or B* is not doxastically justified (since it is not based on a reliable indicator reason). Hence, Smith does not know that *F or B*.

In Kripke County, *blue-barn sensations* are reliable indicators of blue barns. Whenever Henrietta believes that there is a blue barn before her on the basis of a blue-barn sensation, her belief is true. What's important to realize is that, in Kripke County, *blue-barn sensations* are also reliable indicators of barns. Whenever Henrietta believes that there is a barn in front of her on the basis of a *blue-barn sensation*, there is a barn in front of her. Since visual barn sensations don't come colorless, Henrietta's blue-barn sensation is a reliable indicator of both blue barns and barns (in Kripke County), and so, she knows both that there is a blue barn in front of her and that there is a barn in front of her. Epistemic closure is preserved. (Note, however, that *red-barn sensations* are not reliable indicators of barns or red barns in Kripke County. So, if Henrietta were to believe that there is a red barn in front of her on the basis of a red-barn sensation, she would fail to know that there is a red barn in front of her, even if the structure before her were a red barn.)

Internal reasons reliabilism also yields the intuitively correct verdict on both Bon-Jour's Norman case and Lehrer's TrueTemp case. Why? Because both Norman and TrueTemp lack internal reasons for their respective beliefs. Norman just wakes up with the belief that the President is in New York City, and TrueTemp just finds himself with inexplicable temperature beliefs. Since Norman lacks an internal reason for his belief that the president is in New York City [*N*], he is not personally justified in believing that *N* and doesn't know that *N*. Since TrueTemp lacks internal reasons for his belief that the temperature is 104 degrees [*D*], TrueTemp is not personally justified in his belief that *D* and, hence, doesn't know that *D*. Furthermore, since both Norman and TrueTemp lack internal reasons for their respective beliefs, they

---

[19] Strictly speaking, if *that very reason R* is a doxastic reason, then in addition to being a reliable indicator of *p*, *S* must know that *R*, in order for *S*'s belief that *p* to be doxastically justified.

ipso facto lack reliable-indicator reasons, as well. Hence, neither Norman's belief nor TrueTemp's belief is doxastically justified.

What about Hiller and Neta's devastating counterexample to Pritchard's safe true belief account where Smith justifiably infers a safe belief that *F or G* from her justified but false and unsafe belief that *F*? Recall the problem arose because, unbeknownst to Smith, Guinevere possesses absolutely no musical talent. So, there is no nearby world where Guinevere wins a Grammy. But now it looks like Smith's reason *F* for believing *F or G* will be a safe indicator of *F or G* (just because *G* is true in all nearby worlds!). So, doesn't my account also fall prey to their counterexample? The answer is "No." Smith's reason for believing *F or G* is her belief that *F*, and according to condition (3) of DJ, whenever one's reason *R* for *p* is a doxastic reason (i.e., a belief), one must *know* that *R* is true in order for one's belief that *p* to be doxastically justified on the basis of *R*. Since Smith doesn't know that *F* is true (because *F* is false), Smith's belief that *F* does not doxastically justify her belief that *F or G*. True, *she* is personally justified in believing that *F or G* because she bases her belief that *F or G* on her undefeated internal doxastic reason that *F*, which she is personally justified in holding; but her *belief* that *F or G* is not doxastically justified because it is based on a doxastic reason, *F*, that she does not know to be true. Since knowledge requires doxastic justification and since Guinevere's belief that *F or G* is not doxastically justified, it follows that Guinevere doesn't know that *F or G*. In short, DJ rightly precludes cases of knowledge-destroying veritic luck from counting as cases of knowledge.

What about *hole in one*? My belief that not all seventy-two players in the 2022 Masters will shoot a hole in one on Flowering Crab Apple in the first round of play [∼ *H*] is based on a number of undefeated doxastic reasons that I *know* to be true. For example, I *know* that, in the eight-plus decades of play at The Masters, Flowering Crab Apple has yielded only *one* hole in one ever. I also *know* that the odds of a professional golfer shooting a hole in one on any PGA par 3 is rough 1 in 3500. And I *know* that the probability of 72 golfers hitting a hole in one on the same hole in one in single round is 5.608283e-256.[20] Given these ultimately undefeated reasons, I am personally justified in believing that ∼ *H*. Moreover, these reasons are reliable-indicator reasons. There is no nearby world where I have these reasons and every player in the 2022 Masters shoots a hole in one on Flowering Crab Apple in Round One. So, my *belief* that ∼ *H* is doxastically justified. I'm also a reliable interpreter of my reasons for ∼ *H*. So, *internal reasons reliabilism* yields the correct result that I do, indeed, *know* that ∼ *H*.

---

[20] Or equivalently: 5.608283 divided by 10^{256}, which equals a probability of: .00560283. There is no nearby world where I know the above probability *and* every golfer hits a hole-in-one on Flowering Crab Apple in round one of the 2022 Masters. Every nearby world where I have the above reason is a ∼ *H*-world. So, the above reason and the other reasons mentions in the body of the paper are *safe reliable-indicator reasons* for believing ∼ *H*.

11 Conclusion: The IRR Scorecard

Let's see how IRR fares with respect to the desiderata spelled out at the outset. First, IRR resolves the paradox of epistemic luck as follows: we are *evidentially lucky* to live in a veridical world α where many (if not most) of our internal reasons are reliable-indicator reasons, for in a demon world those same internal reasons would not be reliable-indicator reasons.[21] However, it is not a matter of luck when beliefs based on our reliable-indicator reasons in α turn out to be true in α. So, our true beliefs in α based on *safe reliable-indicator reasons* in α are *not veritically lucky* in α. Since knowledge requires doxastic justification (as per K) and since a belief that *p* is doxastically justified only if it is based on an ultimately undefeated internal reliable-indicator reason for *p* (as per DJ), K succeeds in blocking knowledge-destroying veritic luck and, thus, averts the Gettier problem.

Second, IRR's reliabilist account of doxastic justification is immune to the generality problem. The generality problem arises for *process reliabilism* because process reliabilism ties the degree of doxastic justification to the degree of reliability of the process type (instantiated by the process token) that gave rise to the belief. Since any given process token will be an instance of indefinitely many process types, without an account of process relevance that specifies the one relevant process type, process reliabilism is "radically incomplete" (Conee & Feldman, 1998, 2), for without an account of process relevance, process reliabilism has no implications about the justification of beliefs. *Internal reasons reliabilism* is immune to this worry because DJ ties doxastic justification to the *specific internal reason*, *R*, that personally justifies *S* in believing that *p*. Once the internalist has identified the *specific internal reason*, *R*, that personally justifies *S* in believing that *p*, DJ does have clear implications as to whether or not *S*'s belief that *p* is doxastically justified. If *that very internal reason*, *R*, is a reliable indicator of *p* in *W*, then *S*'s belief that *p* is doxastically justified in *W*.[22] If *that very internal reason*, *R*, is not a reliable indicator of *p* in *W*, then *S*'s belief that *p* is not doxastically justified in *W*.

Third, IRR resolves that internalism/externalism controversy by distinguishing two conceptions of epistemic justification—personal justification and doxastic justification—and arguing that internalists employ the former and externalists employ the latter. Recognizing that internalists and externalists have been evaluating the justificatory status of different domains—persons and beliefs, respectively—explains why they have such strong competing intuitions about the very same cases. Once we recognize that internalists have been evaluating the justificatory status of persons and externalists are evaluating the justificatory status of beliefs, we discover that their respective intuitions are correct when it comes to their respective domains of evaluation. As internalists rightly maintain, my demon-deceived *twin* is just as *personally justified* in believing that he is drinking a cup of coffee [*C*] as I am, because we have identical internal reasons—phenomenologically indistinguishable visual, olfactory, and gustatory coffee sensations—for thinking that we are drinking a

---

[21] Fortunately, as we saw in Sec. 4, evidential luck is compatible with knowledge.

[22] Again, strictly speaking, if *that very reason R* is a doxastic reason, then in addition to being a reliable indicator of *p*, *S* must know that *R*, in order for *S*'s belief that *p* to be doxastically justified.

delicious cup of freshly brewed Columbian Supremo. However, externalists are right to insist that my twin's *belief* that *C* is *not doxastically justified* in W_D because *those very visual, olfactory, and gustatory coffee sensations* aren't reliable indicators of a cup of coffee in W_D. In addition, internalists are right to insist that personal justification is necessary for knowledge, and externalists are also right to insist the doxastic justification is needed for knowledge. Indeed, if the world is anything like the way we think it is, then the ultimately undefeated internal reasons that personally justify *us* in our beliefs also doxastically justify those same *beliefs*. The fact that doxastic justification piggybacks on personal justification in veridical worlds is, I suggest, why internalists have often thought that only internalistic justification is needed for knowledge.

Internal reasons reliabilism resolves the paradox of epistemic luck, solves the Gettier problem by ruling out veritic luck, is immune to the generality problem, and provides a coherent ecumenical resolution of the internalism/externalism controversy, all while preserving epistemic closure. All told, that's a pretty impressive scorecard![23]

Declarations

References

Alston, W. (1989). *Epistemic justification: Essays in the theory of knowledge*. Cornell University Press.

Bach, K. (1985). A rationale for reliabilism. *The Monist, 68*(2), 246–263.

BonJour, L. (1998). *In defense of pure reason*. Cambridge University Press.

BonJour, L. (1980). Externalist theories of empirical knowledge. *Midwest Studies in Philosophy, 5*, 53–73.

BonJour, L. (1978). Can empirical knowledge have a foundation? *American Philosophical Quarterly, 15*(1), 1–13.

Chisholm, R. (1964). The ethics of requirement. *American Philosophical Quarterly, 1*, 147–153.

Chisholm, R. (1986). The place of epistemic justification. *Philosophical Topics, 14*(1), 85–92.

Comesaña, J. (2010). Evidentialist reliabilism. *Nous, 44*(4), 571–600.

Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies, 89*(1), 1–29.

DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review, 104*(1), 17–52.

Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy, 49*(1), 1–22.

Engel, M., Jr. (1992a). Is epistemic luck compatible with knowledge? *Southern Journal of Philosophy, 30*(2), 59–75.

Engel, M., Jr. (1992b). Personal and doxastic justification in epistemology. *Philosophical Studies, 67*(2), 133–150.

Gettier, E. (1963). Is justified true belief knowledge? *Analysis, 23*(6), 121–123.

Goldman, A. (1967). A causal theory of knowing. *The Journal of Philosophy, 64*, 355–372.

Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: D. Reidel Publishing Company.

Hiller, A., & Neta, R. (2007). Safety and epistemic luck. *Synthese, 158*(3), 303–313.

Klein, P. (1971). A proposed definition of propositional knowledge. *Journal of Philosophy, 68*, 471–482.

Kung, P. (2010). On having no reason: Dogmatism and Bayesian confirmation. *Synthese, 177*(1), 1–17.

Koppelberg, D. (2012). The significance of disagreement in epistemology. In C. Jäger & W. Löffler (Eds.), *Epistemology: Contexts, values, disagreement, Proceedings of the 34th International Ludwig Wittgenstein Symposium Kirchberg am Wechsel, Austria 2011* (pp. 305–317). Frankfurt: Ontos.

Lehrer, K. (1990). *Theory of knowledge*. Westview Press.

Lehrer, K., & Cohen, S. (1983). Justification, truth, and coherence. *Synthese, 55*(2), 191–207.

Lehrer, K., & Paxson, T., Jr. (1969). Knowledge: Undefeated justified true belief. *The Journal of Philosophy, 66*, 225–237.

Littlejohn, C. (2009). The externalist's demon. *Canadian Journal of Philosophy, 39*, 399–434.

Murphy, P. (2005). Closure failures for safety. *Philosophia, 33*(1), 331–334.

Myers, R., & Stern, K. (1973). Knowledge without paradox. *Journal of Philosophy, 70*, 147–160.

Nozick, R. (1981). Knowledge and skepticism. In R. Nozick (Ed.), *Philosophical explanations* (pp. 167–288). Harvard University Press.

Pollock, J. (1986). *Contemporary theories of knowledge*. Totowa, NJ: Rowman and Littlefield.

Pritchard, D. (2003). Virtue epistemology and epistemic luck. *Metaphilosophy, 34*(1/2), 106–130.

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

Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives, 13, Epistemology*, 141–154.

Sosa, E. (2000). Skepticism and contextualism. *Philosophical Issues, 10, Skepticism*, 1–18.

Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford University Press.

Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy, 65*(6), 157–170.

Vogel, J. (1999). The new relevant alternatives theory. *Philosophical Perspectives, 13, Epistemology*, 155–180.

Williamson, T. (2000a). Scepticism and evidence. *Philosophy and Phenomenological Research, 60*, 613–628.

Williamson, T. (2000b). *Knowledge and its limits*. Oxford University Press.