

## Chapter Seventeen

# A Critique of Prichard's Anti-Luck Virtue Epistemology

**Abstract:** Duncan Prichard is a prolific author. Some of his thinking has evolved over the decades. His preferred species of virtue epistemology is an 'anti-luck' (or 'anti-risk') epistemology. In this chapter we initially examine "Anti-Luck Epistemology" (2007) and *Knowledge* (2009). We will then review Prichard's "The Gettier Problem and Epistemic Luck" (2019) where many of the earlier writings are reiterated and summarized. Prichard has more recently used the term 'epistemic risk' rather than 'epistemic luck,' and argues for an 'anti-risk epistemology' in "Varieties of Epistemic Risk" (2022).

### "Anti-Luck Epistemology" (2007)<sup>1</sup>

In the introduction to "Anti-Luck Epistemology," Prichard says this:

It is a platitude to say that knowledge excludes luck. Indeed, if one can show that an epistemological theory allows 'lucky' knowledge, then that usually suffices to warrant one in straightforwardly rejecting the view. Even despite the prevalence of this intuition, however, very few commentators have explored what it means to say that knowledge is incompatible with luck. In particular, no commentator, so far as I am aware, has offered an account of what luck is, and on this basis identified what it means for a true belief to be non-lucky.<sup>2</sup> In outline at least, this is what I propose to do here. As we will see, this approach casts light on a number of issues in epistemology, and I focus on two issues in particular. First, I will argue that an anti-luck epistemology explains the attraction of adducing a safety condition on knowledge, and also assists us in determining the best formulation of the principle.<sup>3</sup> Second, I will claim that an anti-luck epistemology can show how to best respond to the so-called 'lottery puzzle.'

---

<sup>1</sup> Prichard, Duncan 2007. 'Anti-Luck Epistemology' *Synthese* 158: 277-297.

<sup>2</sup> That there are no prior accounts of what "luck" is, is easily explained. The term 'luck' represents a 'group resemblance concept.' A group resemblance concept is a concept that characterizes kinds of things (i.e., entities, such as 'luck'), where the individual examples (or items) of the kind have an apparent superficial resemblance or loose similarity among instantiations, but there are no necessary and sufficient conditions that theoretically define the kind. Group resemblance concepts have a normal use and if we want, we can construct a unified characterization (or disjunctive definition) of the standard use of these concepts.

<sup>3</sup> *The worth (or value) of this analytical project* to adduce a (single) necessary (e.g., anti-luck, safety, or sensitivity) principle which applies to all instances of 'knowledge' and responds to key epistemological issues (e.g., skepticism, Gettier situations, lottery questions, etc.) *can be questioned*. Outside of prolonged metaphysical discussion of possible worlds, such formulations are mostly idle otherwise. Prichard admits that anti-luck epistemology is not in the business of presenting a complete theory of knowledge.

Prichard asks us to consider a paradigm case of a lucky event: **S** possesses a winning lottery ticket. What is it that makes this event lucky? <sup>4</sup> He says that intuitively, this is an event which obtains in the actual world but doesn't obtain in a wide class of nearby possible worlds (where the typical initial conditions are the same). In *most* nearby possible worlds, that meet this description, the ticket lost. <sup>5</sup>

Prichard then asks us to consider a paradigm example of an event which *isn't lucky*, such as when a skilled archer in good environmental conditions hits his target with an arrow. In this case, not only does the archer hit the target in the actual world, but also in nearly all, if not all, of the nearby worlds (where initial conditions are the same) <sup>6</sup>:

There is no wide class of nearby possible worlds in which the relevant initial conditions for this event are the same in the actual world and yet the event in question does not obtain.

Cases like these suggest the following rough account of luck:

An event is "*lucky*" *iff* it obtains in the actual world but does not obtain in a wide class of nearby possible worlds in which the relevant initial conditions for that event are the same as in the actual world.

---

<sup>4</sup> What makes an event lucky? The following are examples of cases where **S** possesses a true belief, but is in a "lucky" (i.e., accidental, coincidental) circumstance:

- (1) **S**'s true belief **p** ('the answer is c') is the result of *guessing* an answer to a multiple-choice question.
- (2) **S**'s true belief **p** ('that it won't rain') is the result of *wishful thinking*.
- (3) **S**'s true belief **p** ('**S1** doesn't like lamb' is inferred on evidence that she just ordered chicken from the restaurant's menu) is obtained by invalid or fallacious reasoning, or *weak evidence*.
- (4) **S**'s true belief **p** ('It is exactly 12PM') is the result of personally justified belief, but **p** is true for reasons unrelated to its truth (the clock is broken and stopped exactly twelve hours ago).
- (5) **S**'s true belief **p** ('I see a barn') is the result of personally justified belief based on truth-connecting evidence, but **p** is true despite the fact that in **S**'s environmental situation (there are 99 barn facades nearby) it was improbable that **S** had formed a true belief.

**S**'s 'luck' can often be attributed to probabilistic unlikelihood, unreliable methodology, mistaken or errant reasoning, weak evidence, being in a Gettier situation, or being in a Harman situation.

<sup>5</sup> It is striking that Prichard believes that what makes winning a lottery 'lucky' involves intuitions about 'possible worlds.' I believe most persons intuitions involve 'probabilistic unlikelihood.' Di Yang in "What's Wrong with Modal Conceptions of Luck and Risk" (2021) concurs: "Why should we not think that our judgements about modal frequency is inferred from and ultimately reduces to judgments about probabilistic likelihood?" (p .777). ... "it is unclear why we need a modal account of luck at all, especially when probabilistic accounts are more intuitive and easier to apply" (p. 779).

<sup>6</sup> Similar example: In 2019 it was probable (under normal environmental and health conditions) that Green Bay's Mason Crosby would succeed in attempting an extra point (from twenty yards) in an NFL game based upon skill. Crosby's success rate was 97.6% in that year (40 made field goals in 41 attempts).

This definition of “lucky” seems incomplete and the term 'relevant' is dreadfully vague, and Prichard acknowledges both of these problems. (Also, only events which are significant to human agents count as lucky). But Prichard says that this rough account of luck will suffice to get a handle on what an anti-luck epistemology would look like.

Of course, Prichard is not concerned with persons who ‘luckily’ comes across evidence (e.g., by accident, finding a map to a buried treasure) which generates their knowledge. **S** is lucky in finding valued evidence. ‘Luck’ as a manner of how knowledge is obtained, doesn’t undermine knowledge. Luck only undermines knowledge when it impacts directly on the *truth* of the belief in question. With this account of ‘luck’ in mind, a rough specification (stipulative fixed definiens definition) is suggested:

**Lucky True Belief:** **S**'s true belief is lucky *iff* there is a wide class of nearby possible worlds in which **S** continues to believe **p**, and the relevant initial conditions for the formation of that belief are the same as in the actual world, and yet the belief is false.

He then discusses two case studies of what are called examples of “lucky true belief”:

(1) In the first example, **S** actually believes that "Lucky Lass will win the fourth race" by just guessing, and it turns out true that the horse wins. In a wide class of nearby possible worlds (i.e., worlds forming this belief by guess), **S**'s belief is false. "After all, there will be a wide class of relevant nearby possible worlds in which Lucky Lass does not win, but **S** continues to believe by guess that Lucky Lass will win." In other words, in relevant nearby worlds, Lucky Lass doesn't win, and in the actual world **S** possesses a lucky true belief, but not knowledge.

Prichard compares this with where "Lucky Lass will win the fourth race" *can be known* by **S**, if the race was personally ‘fixed’ by a medical assistant **S** by drugging the other horses with a sedative. In nearby worlds, with **S**'s evidence, Lucky Lass wins.

(2) Next, Prichard asks us to consider the standard Gettier-style stopped clock case: **S** looks at the town clock and believes it 12:00 noon, it by happenstance this belief is true, despite the fact that the clock stopped functioning at exactly midnight twelve hours earlier. Prichard says that the Lucky True Belief principle captures the fact that **S** doesn’t know that it is noon (even if it is true) because the belief was lucky.

With the above definition of what a lucky true belief is, and the intuition that knowledge excludes luck, Prichard proposes a necessary condition for knowledge:

**Anti-Luck Condition:** S's true belief is non-lucky *iff* there is *no* wide class of nearby possible worlds in which S continues to believe **p**, and the relevant initial conditions for the formation of that belief are the same as in the actual world, and yet the belief is false.

The anti-luck condition is similar to an epistemic principle defended as a necessary condition for knowledge: the safety principle, a version which is proposed by Ernest Sosa (1999). Roughly, the safety principle can be characterized as follows:

**Safety principle:** S's belief is 'safe' *iff* in most nearby possible worlds in which S continues to form her belief about **p** in the same way as in the actual world the belief continues to be true.

Prichard says that there are a number of advantages to adducing (i.e., offering as argument, reason, or proof) such a principle on knowledge, the chief among them of using safety to show that our denials of skeptical hypotheses (e.g., 'I don't know I'm not a BIV'), while not sensitive to falsehood or justified by accessible grounds, does have "some positive epistemic support in virtue of being safe." Prichard says that we find the 'safety principle' is intuitive *because* it captures our 'anti-luck intuition' about knowledge. This is explicitly expressed in the 'anti-luck condition.' He says:

Since we might reasonably suppose that the 'relevant initial conditions for the formation of the belief' roughly corresponds to the 'way' in which the belief was formed, and that 'in no wide class of nearby possible worlds not-X' is roughly equivalent to 'in most nearby worlds X, then 'safety' and 'anti-luck' are very similar... we find the safety principle intuitive *because* it captures our anti-luck intuition about knowledge which is explicitly expressed in 'anti-luck.'

Prichard says that in upcoming discussions, he'll continue to refer to those 'nearby possible worlds' in which the *relevant initial conditions* for the formation of S's belief are the same as in the actual world and S *continues to form a belief* in the target proposition, as the 'relevant' nearby possible worlds.<sup>7</sup>

---

<sup>7</sup> Prichard stipulates a definition of "relevant" as referring to "nearby possible worlds where initial conditions for the formation of S's belief are the same as in the actual world." This is an example of a formalized abbreviatory fixed definiens definition (i.e., 3b and 3c definiens forms, see chapter 6).

Prichard says that with the 'anti-luck condition' in mind as the underlying motivation for the 'safety principle,' some of the problems raised for safety theories can be addressed. He considers the lottery case (paraphrased):

**The Lottery Paradox:** S doesn't know that **p** (his ticket is a loser), despite excellent statistical grounds for believing that it is a loser.

This case is about S's *strong belief*, prior to the draw being announced, that 'he owns a *losing* ticket.' Intuitively this is *not* a case of knowledge, even despite the excellent statistical grounds S has in support of the belief.

Prichard says this happens because when the belief is true, it is nevertheless a matter of luck that the belief is true.<sup>8</sup>

After all, there are relevant nearby worlds in which S's belief is false (i.e., those worlds in which S owns the winning lottery ticket).

If this is right, however, then it seems that demanding S's belief be true in 'most' nearby possible worlds, as the safety principle demands, will not suffice, since, intuitively, this condition is met in the lottery case where there are only very few relevant nearby worlds in which one forms a false belief in the target proposition. The safety principle, thus, predicts knowledge in this case, even though our intuitions, guided by the anti-luck platitude, dictate otherwise.

Of course, the safety principle is only meant to be a necessary condition on knowledge. It could be argued that there are further conditions on knowledge available to deal with this case, but since the intuition here is that knowledge is lacking *because* of the presence of luck, and since we have noted, safety turns out to be a natural way of understanding the anti-luck condition, this dialectical move is not very appealing. Accordingly, it seems that one must understand safety along much stronger lines in demanding that S's belief be true not just in most of the relevant nearby worlds, but in nearly all (if not all) of them.

We now obtain a stronger safety principle\*:

**Safety principle\*:** S's belief is 'safe' iff in nearly all (if not all) nearby possible worlds in which S continues to form her belief about the target **p** in the same way as in the actual world the belief continues to be true.

---

<sup>8</sup> Prichard's intuition is that it is a matter of 'luck' that S possesses a true belief that her ticket is a loser. This 'luck' explains why S doesn't know that her ticket *x* is a loser (prior to the drawing). Contrast the PE explanation: S doesn't know that ticket *x* is a loser (prior to drawing) because she cannot possess 'relevant reasons' (violating PE condition 3). The concepts of 'luck' and 'relevant reasons' are here contrasted.

However, strengthening the safety principle may make it unable to deal with mundane examples of knowledge possession. He considers this case study from Sosa (2000, 13):

**The Rubbish Chute Case:** On my way to the elevator, I release a trash bag down the chute from my high-rise condo. Presumably, I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) was not to arrive there? That presumably would be because it had snagged somehow in the chute on the way down (an incredibly rare occurrence), or some happenstance.

Do we want to say that Sosa has knowledge that his bag is in the basement? Intuitively, we do, but if safety is strengthened along very robust lines as SP\*, then it has been claimed, this seems to be ruled-out since surely there are quite a few nearby possible worlds, those in which the bag snags, in which Sosa continues to form his belief on the same basis and yet forms a false belief as a result. We seem to be stuck between two opposing intuitions:

- (1) Safety (SP\*) must be understood robustly (i.e., strong) in order to eliminate the luck in play in the lottery case.
- (2) Safety (SP) must be understood as weaker, in order to accommodate mundane cases of knowledge such as the rubbish chute example.

Closer inspection of the rubbish chute example reveals, however, that the challenge it poses to SP\* is not nearly as clear as many have thought. Sosa says that the snagging of the of the bag would be "an incredibly rare occurrence." It seems that Sosa's belief *isn't* lucky since the probability that the bag snags is quite remote... On this reading of the example, the intuition that Sosa knows that his bag is in the basement is not under threat, but now it isn't so obvious that there are very many (if any) relevant nearby possible worlds in which his belief is false—the worlds in which the bag snags on the chute seem relatively far off.

In contrast, we can stipulate the details of the case such that it is clear that Sosa's belief *is* lucky by demanding that there is a wide class of relevant nearby worlds in which S forms a false belief because the bag has snagged in the chute. In order to make this supposition plausible we would have to imagine, for example, that there is a snag in the chute that the bag is *almost* snagging each time. If that's right, however, then it would be odd to think that Sosa does know that his bag is down in the basement, since it is clearly a matter of luck that his belief is true in this regard given the nearness of the relevant error-possibility.

Everything thus depends on how we are understanding the details of the case. Filling in the details here thus highlights that formulating safety in a fairly strict fashion as SP\* so as to deal with the lottery case may not result in our denying knowledge in cases like the rubbish chute example after all.<sup>9</sup>

After several paragraphs of ruminations about 'luck' and 'possible worlds,' Prichard presents a new formulation of 'safety':

**Safety principle\*\*:** S's belief is 'safe' iff in most nearby possible worlds in which S continues to form her belief about the target **p** in the same way as in the actual world, and in all very close nearby worlds in which S continues to form her belief about the target **p** in the same way as in the actual world, the belief continues to be true. (Another fixed definiens stipulative definition).

He says:

With SP\*\* in mind, consider again the formulation of the rubbish chute example in which it clearly isn't lucky that Sosa's belief that his rubbish is in the basement is true (the formulation in which the possibility of the bag snagging is actually quite remote).

This formulation was such that although one might plausibly suppose that there were some relevant nearby possible worlds in which Sosa formed this belief and yet his belief was false, it was certainly the case that (unlike in the lottery example) none of the *very close* possible worlds fitted this description. Paying close attention to the way in which our judgements about luck are affected by the closeness of the relevant nearby possible worlds thus highlights a relatively minor modification that we need to make to our formulation of the safety principle in order to meet this particular problem.

---

<sup>9</sup> Does S know **p**: 'My bag is in basement?' and if so, why, and if not, why not? With the PE definition of knowledge, like the parked car case, whether S knows **p** is in part determined by one's material situation. We are inclined to say that you know that your bag is in the basement, given that when releasing the bag, the chute normally functions in a way for the bag to reach the basement. But what about the possibility of a rare snag? For S to know **p**, all four conditions of the PE definition must be satisfied and so typically S knows **p**. This is the normal situation for a well-maintained high-rise (that Sosa assumes). But, if conditions 1, 3, and 4b are false (e.g., the bag is snagged in the chute, after normal wear, or bad installation, or squirrel nest), then S *fails* to know **p**. If S is (somehow) aware of the live possibility of a malfunction, and for some reason cannot dismiss this counter-possibility **q**, then S doesn't know **p**. In a case of S's knowledge (i.e., the trash *is* in the basement), S fallibly discards undermining possibilities (e.g., a snag), and there *just are* (i.e., *exist*) *no other unconsidered facts* (i.e., *true propositions*) that defeat (or substantially undermine) S's reasons for believing **p**. In determining whether 'S knows **p**,' the contingencies of S's external environment take precedence. This answer is the standard to similar cases: (1) Who is the current US President? (Perhaps a sudden fatal heart attack has transferred office to the Vice President). (2) Did my football team win? (There might have been a misprint when reading a score from the newspaper).

Prichard then considers Hawthorne's (2004) lottery problem:

- (1) S knows that he won't have enough money to go on safari this week.
- (2) If S knows that he won't have enough money, then he is in a position to know he won't win a major lottery prize this week. (Epistemic closure).
- (3) Conclusion: S is in position to know that he won't win a major lottery prize this week.

Prichard believes that clarifying the role of luck in this example will show why there is, in fact, no puzzle here.

Prichard responds to the puzzle by saying that there is an ambiguity in how we are reading (1) and (3). If we understand the facts of the situation in such a way *that (1) is intuitive*, then our *intuitions that (3) is false subside*.<sup>10</sup>

Conversely, provided that we consistently understand the facts of the situation in a way that *(3) is counterintuitive* (because there are nearby possible worlds in which S wins a prize but continues to believe that he hasn't) then we no longer find (1) intuitive.<sup>11</sup>

He closes this section with the following paragraph:

The lottery argument thus plays on an ambiguity in how we are reading (1) and (3). Provided that we consistently understand the facts of the situation in such a way that (1) is intuitive, then our intuition that (3) is false subsides. Conversely, provided that we consistently understand the facts of the situation in such a way that (3) is counterintuitive, we no longer find (1) intuitive. Taking anti-luck epistemology seriously thus enjoins us to be crystal clear about what the facts of the situation are in cases like this and to keep them fixed throughout. Once we do that, however, there is no tension between the two claims, and thus there is no lottery problem that needs to be responded to. Anti-luck epistemology is thus able to resolve a key puzzle in contemporary epistemology, and thereby undermine the motivation that many have seen in this puzzle for a form of epistemological revisionism.

---

<sup>10</sup> This doesn't seem right. Our intuitions that (3) is false should not subside.

<sup>11</sup> In chapter 1, it is argued that both premises in Hawthorne's argument are false, as well as the conclusion. This paradox is explained by recognizing that there are *two* ordinary language senses of 'relevant.' In one sense, S *cannot* have relevant (truth-connecting) evidence to *know* she won't win the lottery and become rich, but in another sense, S has 'relevant' (wider, related) probabilistic evidence that there is little reason (low wages, no inheritance, lottery win improbable) to believe that winning the lottery would become true. S's claim to *know* that he'll *never* afford to go on safari is *false*, (premise 1) given a purchased lottery ticket.



### ***Knowledge (2009)***

With *Knowledge (2009)*, Prichard offers a perspective and interpretation of knowledge presented in an introductory textbook form. In this book and in the following essays, there is some repetitive overlap of ideas. Here he presents two worries about the 'analytic project' of epistemology (p. 3):

(1) Is knowledge the kind of thing that one can analyze?

(2) What does it mean to offer an analysis of knowledge?<sup>12</sup>

With respect to the second question, one way to offer an 'analysis' is to provide the necessary and sufficient conditions that *capture our everyday use(s) of the term 'knowledge'* that allows us a belief (a judgment) about whether (intuitively) **S** has knowledge of **p** in various case studies. Prichard says that this traditional analytic project is *hopeless* (italics added):

For one thing, there is very little reason to think our *folk usage* is going to be disciplined enough such that it generates an extension that a particular set of necessary and sufficient conditions could capture.<sup>13</sup>

But we don't want an analysis of the term 'knowledge' that is completely theoretical and divorced from everyday folk usage. He says:

What we are after, then, is an analysis *anchored* in our everyday usage, but which is suitably *'cleaned up'* and is in this sense the product of theory. Of course, this doesn't mean that we can summarily dismiss any recalcitrant linguistic data uttered by the folk, since it is incumbent upon us to explain away this data in a compelling manner where applicable. The point is just that we should not think that our project is hostage to this data in the way that the overly austere conception of the analytic project suggests (p. 4, italics added).<sup>14</sup>

---

<sup>12</sup> In treating knowledge as a natural kind entity, the PE definition seeks the necessary and sufficient conditions that captures (an extension for) *when S knows p* (and when S doesn't know p) and *why*.

<sup>13</sup> Although he PE definition seeks the necessary and sufficient conditions that captures (an extension for) *when S knows p* (and when S doesn't know p) and *why*, its analysis doesn't focus on 'ordinary folk usage' of a term as much as 'case studies' that critically analyze our implicit epistemology-related beliefs. (Also, talk of 'extensions' and 'sets' shouldn't be involved in analyzing what 'knowledge' is).

<sup>14</sup> Prichard uses analogies: An analysis "anchored" in everyday use but suitably "cleaned-up" in leading the way to a possible worlds model with its multitude of complex theoretical stipulative definitions.

Prichard is offering:

(1) An analysis which captures what is central to our folk usage—i.e., the paradigmatic use of this term.<sup>15</sup>

(2) A modest theoretical ambition when compared with the austere reading of the analytic project. He says that with an austere rendering of the analytic project, it becomes very depressing because it is easy to find counterexamples to necessary and sufficient conditions definitions. Once one moves away from an austere reading, then such counterexamples cease to be so readily available.

After presenting intuitions from two case studies about ‘Gullible John’ and ‘Scientist John,’ that introduce a *doxastic* condition, a *factivity* condition, and the notion of ‘luck,’

Prichard suggests these two intuitions (p. 9):

(1) When **S** has knowledge, **S**'s true belief could not easily have been wrong. Call this the *anti-luck intuition* about knowledge.

(2) Knowledge is the product of cognitive ability. **S** must acquire a true belief through cognitive ability. Call this the *ability intuition* about knowledge.<sup>16</sup>

Prichard presents one of Gettier's original examples (about Smith having ten coins) and an example from Chisholm (sheep in a field) as well as ‘Henry and the Barn.’<sup>17</sup> Prichard says (paraphrased, italics added, p. 13):

While **S** *isn't* looking at a sheep, and Henry *really* sees a barn, this means that it isn't even essential to Gettier-style cases that **S** is in error in how a belief is formed; *all that is essential* is that the *justified true belief* in question is only *true as a matter of luck* (and thus the anti-luck intuition is compromised).<sup>18</sup>

---

<sup>15</sup> (The PE theory is not primarily concerned with the paradigm *use* of the ‘*term*’).

<sup>16</sup> I share *neither* of these ‘worldview’ intuitions, especially the ‘ability condition.’

<sup>17</sup> Two kinds of luck can be distinguished (as follows):

(1) Gettier situation: That **S**'s belief is true, has *nothing* to do with **S**'s reasons for believing **p**. **S** has a personally justified true belief that is only true by “luck” (1,2, 4a satisfied, with 3 and 4b violated).

(2) Harman situation: That **S**'s belief is true, has *something* to do with **S**'s reasons for believing **p**, but there are true propositions (violating 4b) that if accessed by **S**, might undermine **S**'s justification (4a) and strength of belief (2), which *might* lead **S** to not know **p**. **S** has a personally justified true belief that is only true by “luck” (1, 2, 3, 4a are satisfied, but 4b is violated).

<sup>18</sup> To reiterate, in consolidating Gettier cases and Harman cases under the concept of ‘luck,’ Prichard explicitly seeks one or two necessary principles for knowledge that dictate whether **S** knows or doesn't know when addressing key epistemological questions and exemplary case studies. To repeat, according to Prichard, the traditional necessary and sufficient conditions analysis of knowledge is “hopeless.”

At pp. 14-15, Prichard asks, "... how should one go about formulating a theory of knowledge that is Gettier-proof? This is the *Gettier problem*."<sup>19</sup>

Prichard says that early responses to the Gettier problem assumed that another condition needed to be added to the JTB account. One suggested condition was that there be 'no false premises' in S's justification for believing that **p**. But Prichard argues that this condition fails because knowledge is consistent with **S** possessing *some false* assumptions, provided that they are relatively minor.

At pp. 15-17, Prichard introduces a distinction between internal and external epistemic conditions, and notes that the Gettier style cases require us to at least *add* some *external* epistemic condition to what knowledge is. He says that the prospects for any theory of knowledge which only has internal epistemic conditions are dim. He then presents a standard drawing of the distinction:

(1) An 'internalist theory' of knowledge is any theory of knowledge which incorporates at least one internal epistemic condition.

(2) In contrast, an 'externalist theory' of knowledge is any theory of knowledge which *isn't* an internalist theory of knowledge—i.e., which doesn't insist on there being an internal epistemic condition.

One might respond to the Gettier problem by just adding an 'externalist condition' to 'true belief' to produce knowledge. Process reliabilism is an example:

**Process reliabilism (df.):** 'Knowledge' is true belief that is the product of a *reliable process*, where a "reliable process" is a process that tends to result in true beliefs.<sup>20</sup>

---

<sup>19</sup> Should we be '*formulating*' a theory of knowledge which is *Gettier proof*? Prichard believes that any response to the Gettier problem *must include anti-Gettier condition(s)*. Is this the Gettier Problem? (It's not my Gettier problem). Instead, an *explanation* of Gettier and Harman cases is needed. In short, a *description* of these kinds of cases (again in greater detail) are as follows:

(1) In Gettier cases, S is personally justified in believing p, but S's *reasons* for believing p are *irrelevant* to the truth of p. There exists evidence (not accessed by S) that would weaken (i.e., undermine, undercut) S's belief that p. The truth of p is just coincidentally (i.e., accidentally) true, given S's *evidence*. S is just 'lucky' to have a JTB that p.

(2) In Harman cases, S is personally justified in believing p, and S's *reasons* for believing p are *relevant* to the truth of p, but there exists evidence (not accessed by S) that would weaken (i.e., undermine, undercut) S's belief that p. The truth of p is just coincidentally (i.e., accidentally) true, given S's *environment*. S is just 'lucky' to have a JTB that p.

The epistemic condition laid down by the reliabilist is clearly an external epistemic condition since **S** cannot know by reflection alone that his belief is formed in reliable fashion (**S** can know by reflection with good reason to think that a belief is formed in a reliable fashion, but that's a different matter entirely).

Prichard discusses the 'chicken sexer' example (p. 19):

**S** has a highly reliable ability to distinguish between male and female chicks. **S** doesn't know how he does this—he thinks it might be through sight and touch, but he's not sure—and he also hasn't taken the trouble to verify that his ability really is reliable (he just takes it for granted that it is). Still, if one wants to know whether a chick is a male or a female, then go to **S** and he'll tell you. Does **S** have knowledge?<sup>21</sup> Here is where externalists and internalists about knowledge diverge. In favor of externalism, notice that the true beliefs formed by **S** respect two core intuitions. On one hand, the beliefs are clearly a product of *genuine cognitive ability*, since it is stipulated that **S** does indeed have the ability to tell reliably male and female chick apart. Thus, the true beliefs in question don't contravene the *ability intuition* about knowledge. Moreover, the true beliefs formed by **S** clearly *aren't true as a matter of luck*, given that they are indeed *formed as a result of genuine cognitive ability*. Thus, the true beliefs in question don't contravene the anti-luck intuition about knowledge either. There are, then, some compelling grounds for supposing that the true beliefs formed by **S** should qualify as knowledge.

In opposition to this, internalists about knowledge insist that **S**'s merely forming a belief in a way that is in fact reliable shouldn't suffice for knowledge, for what is required in addition is some reflective grasp of his epistemic situation. The internalist argues that externalism unduly allows knowledge to be completely opaque to **S** by enabling **S** to have knowledge even in the absence of supporting reflectively accessible grounds.

---

<sup>20</sup> The 'process reliabilism' thesis seems entirely vague and not immediately intuitive.

<sup>21</sup> Let us add a hypothetical contingent fact: Researchers have found out empirically that chicken sexers, like **S**, reliably discriminate based upon *smelling* the chickens. They have causally based true beliefs based upon *olfactory* nerves without perceptual seeing (eyes) or touching (tactile). With the PE definition, and if this new research were true, condition 3 is satisfied. **S** believes a chicken's sex upon some vague reasons (i.e., sight, smell, or some sensory reasons) that are relevant (i.e., truth-connecting) for why **p** should be believed. **S** has *animal knowledge* (satisfying the basic PE conditions 1-4) but it can be debated whether **S** has *human knowledge*. Condition 4a may (or may not) be violated. If the evidence of **S**'s own success in chicken sexing results are good enough for him (to be sufficient evidence) then **S** has human knowledge. But for some internalists, conditions 4a and 4b are violated. There exists evidence **q**, establishing that neither sight nor touch are evidential reasons for **S**'s belief. Although **S** possesses reliably produced true beliefs, **S**'s stated reasons (for believing **x** is male or female) are too imprecise (or slightly in error), so it might be prescribed that **S** doesn't have a sufficient personal justification for allowing that **S** has *human knowledge*.

Prichard asks, how should one formulate the anti-luck condition? Two main proposals are found in the literature. The first (p. 24):

**The Sensitivity Principle:** If **S** knows that **p**, then **S**'s true belief that **p** is such that, had **p** been false, **S** would not have believed **p**.

The basic idea behind the sensitivity principle is that when it comes to knowledge, we don't simply want a belief that matches up with the facts- i.e., which is true—but also a belief that is sensitive to the facts, such that one wouldn't have believed what one did, had it been false.<sup>22</sup>

He returns to consider Chisholm's case study (p. 24):

Mistaking a dog for a sheep, **S** justifiably believes from the way things look to him that **p**: 'There is a sheep in the field.' But **p** is true only because of the presence of an unseen, unsuspected sheep. (Chisholm 1977, p. 105).

Prichard says that we can explain why **S** doesn't have knowledge in terms of the sensitivity principle since **S**'s belief is clearly insensitive. If there hadn't been a hidden sheep, **S** would have continued to believe that there was a sheep (and so formed a false belief).

The proponents of sensitivity have a sophisticated story to tell, which includes the notion of possible worlds. Prichard says that 'possible worlds' are a "useful" philosophical device when it comes to thinking about possible states of affairs. This is to imagine a possible world. We can 'order' possible worlds in terms of *how similar* they are to the actual world—i.e., in terms of *how much* is *different* from the actual world.

He compares himself being in hypothetical situations: (1) Upstairs typing a paper or (2) Downstairs cooking dinner. These are possible worlds and then there is the actual world. There is very little difference if this is all that is different between worlds. This second possible world is very much like the actual world, in that little would be needed to change about the actual world in order to turn it into this possible world. Some possible worlds are very different, such as where the foundational laws of physics are different.

---

<sup>22</sup> The 'sensitivity' principle doesn't strike me as 'intuitive.' It seems narrow and artificial.

We can 'order' possible worlds in terms of how similar they are to the actual world. He says (p. 26):

Possible worlds talk is very useful philosophically, which is why so many philosophers employ it. That such talk is useful does not, however, make it legitimate, and there are some non-trivial worries about possible worlds. For example, when a statement about the actual world is true we have a fairly clear grip on what it is that makes it true—i.e., that it is some feature of the actual world. Clearly, though we cannot straightforwardly apply this reasoning to statements involving possible worlds, since (intuitively at any rate) such worlds don't really exist. There are also worries about the objectivity of possible worlds talk. After all, similarity is very vague, and possibly even context-sensitive notion, and so one might wonder whether there could not be a lot of variability in the truth values that we intuitively attribute to statements about possible worlds. And there are other problems too.

Despite these difficulties, however, I think that we can legitimately employ this framework without further concern for our purposes here. For one thing, it is often the case that when pursuing some area of philosophy, one has to take as given answers which are unresolved in another area of philosophy (for example, a lot of discussion in ethics presupposes the possibility of free will, and yet this metaphysical issue is far from settled).

Accordingly, it is not unusual that as epistemologists we have to help ourselves to a philosophical framework that is philosophically contentious outside of epistemology. However, even if it were to turn out that possible worlds talk is unsustainable, this would not mean the end of the kind of epistemological approach sketched here. Instead, it would simply mean that we would have to recast this approach within a different framework.

He returns to the sensitivity principle (p. 27). It is a 'modal principle,' a principle which implicitly appeals to modal notions like that of a possible world. Using the possible worlds framework, we can be more precise about which possible world is the relevant one to consider when we are evaluating whether a belief is sensitive. The possible world that we are interested in is the one in which everything is the same except that which would need to be different for **p** to be false (i.e., the closest world in which **p** is false).

Prichard says that using 'possible worlds' and 'sensitivity' in the analysis of the sheep in the field example, allows an elegant solution for explaining why **S** doesn't have knowledge. In order to evaluate whether **S**'s belief is sensitive, we need to consider the closest possible world is false—i.e., the world where little else changes other than the

truth of **p**. Such a world would be a world where everything is the same except that there are no sheep in the field. What does **S** believe in this world? In the sheep in the field case, **S** will continue to believe that there is a sheep in the field regardless, since he will still be looking at the sheep-shaped object. In contrast if **S** formed his belief by actually looking at a sheep, then this problem wouldn't arise. The nearest possible world in which **p** is false in this case would be a world in which the sheep that **S** is looking at is no longer there, but everything else is the same. But if there is no sheep there, then **S** wouldn't believe that there is a sheep in the field and so **S**'s belief would in this case be sensitive.

Prichard says that anti-luck epistemology is any epistemology which motivates an epistemic condition on knowledge in terms of its ability to accommodate the 'anti-luck intuition'. More specifically:

... one might motivate such a proposal by explicitly analyzing luck and the sense in which *bona fide* knowledge excludes luck as a means of casting light on the nature of knowledge (p. 38).

This approach differs from the proponents of safety or sensitivity principles that typically take 'luck' as a primitive notion. Despite its importance, little has been written by philosophers about luck. But this is starting to change with a recent proposal that:

In essence, a lucky event is an event that obtains in the actual world but does not obtain in a wide class of near-by possible worlds.

So, for example, that the sun rose this morning is not a lucky event, since it obtains not just in the actual world but also in all near-by possible worlds. Winning the lottery, in contrast, is a lucky event since while it happens in the actual world there are lots of near-by possible worlds where this event doesn't occur (p. 39).

He examines two complexities (paraphrased):

(1) We must restrict the class of relevant worlds to those where the initial conditions for the target event are the same as in the actual world

(2) Luck comes in degrees depending on how close the near-by world is in which the target event doesn't obtain. To illustrate: A sniper misses your head by a few inches. You are lucky to be alive. We can account for this view by recognizing that while you are alive in the actual world, there is a wide class of near-by possible worlds in which the relevant initial conditions for the target event obtain (e.g., someone takes a shot at you) and you are dead just now.

Suppose, however, that the sniper missed by more than a few inches (e.g., a foot). Intuitively, it is still a matter of luck that you are alive, and the account of luck can accommodate this intuition since there is a wide class of near-by possible worlds in which the relevant initial conditions for the target obtain and in which you are killed by the bullet. Nevertheless, intuitively, you are *luckier* to be alive in the former case than in the latter case, and we can account for that in terms of this analysis of luck. For what differentiates these two lucky events is that in the luckier of the two events the near-by possible world in which one is killed is closer. More generally, the closer the near-by world in which the target event obtains, the luckier the event will be.<sup>23</sup>

Prichard asks what sort of anti-luck epistemology does this account of luck suggest? We want to capture the sense in which when S has knowledge, S has a true belief that is not true as a matter of luck: It needs to be the case that S's belief is *true in the actual world*, and yet *false* in a wide range of *near-by possible worlds*. But this is essentially the 'safety principle,' since it demands that S's belief not just be true in the actual world but also in all, or at least nearly all, near-by possible worlds.

Prichard goes on for several pages discussing complexities involving modal principles (sensitivity, safety) that need to be relativized to a method and a strong or weak reading of the safety principle. He says that the general idea behind the safety principle—indeed, behind anti-luck epistemology more generally—is that knowledge entails a true belief that could not have easily been false (p. 45). He concludes that the problem with robust anti-luck epistemology is that it can't completely accommodate the ability intuition (p. 50). We need to turn elsewhere, in chapter three, for a theory to accommodate the ability intuition.

Chapter 3 is entitled 'Virtue Epistemology.' Prichard says that in chapter one, process reliabilism was identified as the theory that 'knowledge' is true belief that is formed with a reliable process, where a reliable process is one that tends to lead to true beliefs rather than false beliefs. He argues that process reliabilism cannot deal with Gettier-style cases and cannot be an adequate account of knowledge as it stands.

---

<sup>23</sup> It seems to me that these are just frivolous intuitions about possible worlds.



He says that Gettier-style cases essentially trade on the anti-luck intuition. Knowledge involves a belief that is true, not simply as a matter of luck. Prichard claims that process reliabilism cannot accommodate this intuition. Process reliabilism is unable to capture the sense in which genuine knowledge reflects a responsiveness (on S's part) to the facts (i.e., evidence). S responds to evidence in an efficient and rational way. Process reliabilism is unable to accommodate the ability intuition (i.e., that the truth of a belief should be a result of cognitive agency).<sup>24</sup>

Prichard discusses a turn from process reliabilism to virtue epistemology. While process reliabilism faced problems because it can't deal with Gettier-style cases and accommodate the ability intuition, a new form of reliabilism was proposed which could, it seems, at least accommodate the ability intuition. He says (paraphrased, p. 53):

This view was known as virtue reliabilism, and it is a form of virtue epistemology. What is common to all virtue epistemological views is that they define knowledge in terms of S's reliable cognitive abilities, such as faculties (e.g., senses) and intellectual virtues (e.g., conscientiousness when dealing with evidence). Just as virtue theories in ethics focus on the morally virtuous S—the S who has the right mix of virtuous moral traits, and so should be admired and emulated—so virtue epistemic views focus on the mix of cognitive traits that S should possess in order to be a 'good' epistemic subject. Virtue reliabilism is one variant of this view which essentially remodels simple process reliabilism along virtue epistemic lines. Here is how it is usually formulated (p. 54):

**Virtue reliabilism:** S knows that **p** iff S's true belief that **p** is the product of the reliable cognitive traits that make up S's cognitive character.

Prichard goes on to discuss a 'modest virtue epistemology' (p. 61):

S knows that **p** iff S's true non-Gettierized belief that **p** is the product of reliable cognitive traits that make up S's cognitive character.

---

<sup>24</sup> Prichard presents a problem for process reliabilism (p. 49). The Temp case study: S's personally justified true belief of the temperature on the thermometer is not the result of the functioning thermometer, because its reading is actively fluctuating randomly within a given range. However, as the result of an unknown intermediary (e.g., a hidden person) who manually (mechanically) adjusts the temperature in the room (quickly and magically) so that it matches whatever the random temperature appears on the thermometer. S forms a belief in a reliable fashion, but lacks knowledge because, in part, the reliability in question does not reflect responsiveness on S's part to the facts. The reason(s) why S's belief is true and reliably formed has little to do with S's cognitive traits. It isn't true of S that his true belief is a product of his reliable cognitive abilities. Instead, it is the product of environmental factors that are entirely independent of S's cognitive traits. Prichard (2019) repeats this example as the case of the 'helpful demon' (p. 105).

By being *less ambitious*, modest *virtue epistemology* can *sidestep* the problem posed by Gettier cases (by *ignoring* them). But the Henry and Barn case is still a tough problem for robust virtue epistemology. Prichard concludes that it appears virtue epistemology is unable to offer us an account of knowledge which can deal with all kinds of knowledge-undermining luck (p. 70).

Chapter 4 is entitled 'Anti-Luck Virtue Epistemology.' In chapter 1, when considering how to analyze knowledge, there were two overarching intuitions:

- (1) Knowledge is due to cognitive ability (the ability intuition).
- (2) Knowledge excludes luck (the anti-luck intuition).

What does it take to ensure one's true belief is not due to luck? (The anti-luck intuition is dominant). Knowledge is essentially non-lucky true belief. But this was also unsatisfactory. There are cases where **S** meets this anti-luck intuition but doesn't know. True belief so formed is not the product of **S**'s cognitive ability. In short, a robust anti-luck epistemology cannot accommodate the ability intuition (pp. 72-73).

Since luck is a modal notion (as discussed in chapter 2), an anti-luck condition will inevitably involve a modal epistemic principle like safety or sensitivity. He introduces the following (paraphrased, p. 74):

**Anti-luck virtue epistemology:** **S** knows that **p** if and only if (i) **S**'s true belief that **p** is the product of the reliable cognitive traits that make up **S**'s cognitive character (such that **S**'s cognitive success is to a significant degree credible to **S**'s cognitive agency), and (ii) **S**'s belief that **p** is safe.

Prichard says that the Gettier cases (and the non-standard ones like Henry and the Barn) are dealt with by the anti-luck condition, since all of these cases have a true belief which is unsafe. The temp case, for why doesn't know the temperature is discussed in above footnote, is dealt with by the ability condition.

With an anti-luck epistemology, a key point is that **S** can form a true belief through cognitive ability in the relevant sense without it being the case that **S**'s cognitive character is the most salient feature in a causal explanation of that cognitive success; instead, all that is required is that **S**'s cognitive agency play a significant part in the acquisition of the true belief (p. 76). In addition, it is assumed that **S** is in an epistemically friendly environment such that **S**'s belief meets the anti-luck condition and hence is safe.

We now turn to chapter 6 entitled 'Radical Skepticism.' Radical skepticism is the view that knowledge (or most of it) is impossible. The thesis about the impossibility of knowledge isn't addressed to real-life skeptics (if there are any), but rather it is a *paradox* that exposes a deep tension among our epistemological concepts. Prichard says that since *paradoxes* are *generated by our own concepts*—in this case our epistemological concepts—it clearly doesn't matter whether there is anyone out there who actually argues for the skeptical conclusion. The radical skeptical claim is that it is impossible for **S** to possess knowledge, which is disturbing. The standard problem is presented involving:

- 1) The Skeptical Hypothesis
- 2) The Closure Principle
- 3) The Denial of Ordinary Knowledge.

Prichard says with respect to any wild skeptical hypotheses, our tendency is to disregard them. If there is no specific reason for believing that a certain error-possibility might obtain, our normal practice is just to ignore the possibilities. Prichard says (paraphrased):

The general thought in play here is what is known as the '*relevant alternatives intuition*' about knowledge: that in order to have knowledge, **S** needs only to be able to rule out relevant error possibilities (i.e., alternatives), and not also the irrelevant ones. And what makes an alternative relevant: Well, intuitively at least, it is whether there is any specific reason for thinking that it the kind of error possibility that could well obtain. Three principles are specified:

**(1) Relevant alternatives principle:** In order for **S** to know **p**, **S** only needs to rule-out relevant error possibilities (i.e., alternatives) and not also the irrelevant ones. What makes an alternative relevant? Whether (or not) **S** has a specific reason for believing that a certain error possibility might obtain.

**(2) Infallibilism:** All error possibilities must be ruled-out before **S** possesses knowledge. This view is contentious. It is in conflict with the 'relevant alternatives' principle.

**(3) Epistemic closure:** If **S** knows **p**, and knows that **p** entails **q**, then **S** knows that **q**. Knowledge is 'closed' under known entailment.

**The closure-based skeptical argument is sketched:**

- 1) It is impossible to know the denials of skeptical hypotheses.
- 2) If we are to know ordinary propositions, we must know the denials of (at least some) skeptical hypotheses.
- 3) Conclusion: We don't have knowledge; it is impossible.

Prichard says that the epistemic closure principle might be the weakest point in this argument and that premise 2 could be rejected. He says that this is easier said than done. He says that the most influential way of objecting to the closure principle is by appealing to the sensitivity principle (p. 109).

Prichard explores why adopting the sensitivity principle might be grounds for rejecting closure. He states (p. 109):

Stated informally, the sensitivity principle demands that one's true belief be such that, in the nearest possible worlds in which what one believes is no longer true, one no longer believes the target proposition. So, for example, my true belief that I have two hands is sensitive because, in the nearest possible worlds in which I don't have hands—where I am staring incredulously at stumps at the ends of my arms right now—I no longer believe that I have hands... In general, for most of the everyday propositions which we believe it is fairly easy to have a sensitive belief in these propositions. Hence, a sensitivity-based epistemology will be inclined to credit us with knowledge of these everyday propositions.

He notes that our denials of skeptical hypotheses are by definition 'insensitive.'

Right now, I believe that I am not a BIV, and this belief is, let us grant, true. Crucially, however, in the nearest possible worlds in which is no longer true that I am not a BIV (i.e., the worlds in which I am a BIV) I will continue to believe that I'm not a BIV regardless since my experiences are indistinguishable from my current experiences. My belief in this proposition is thus insensitive, in that the relevant fact can change, but my belief wouldn't change with it.

One advantage of the sensitivity principle is thus that it can account for why we regard the first principle of the skeptical argument so compelling. According to sensitivity-based epistemology, that is, one lacks knowledge of the denials of skeptical hypotheses precisely because it is impossible to have sensitive belief.

Proponents of the sensitivity principle argue that we can have knowledge of everyday propositions even while lacking knowledge of skeptical hypotheses, and hence epistemic closure must be rejected.

Prichard says that given the plausibility of the closure principle, it is incumbent upon anyone who accepts this sensitivity view to explain why epistemic closure seems true; but is yet false (p. 111).<sup>25</sup> He says that even if we grant that there are such explanations ought we to find this approach to the skeptical problem feasible? Prichard doesn't find this sensitivity and anti-closure approach feasible. In chapter two, it was

---

<sup>25</sup> It is explained in chapter 2 why epistemic closure is false.

shown that a safety principle is more preferable as an anti-luck principle, and it isn't in conflict with closure. He says (paraphrased, p. 111):

(1) Consider first **S**'s belief that **S** has two hands. In order for this to be safe, **S** needs it to be the case that in all near-by possible worlds—certainly most nearby possible worlds at any rate, and all very close near-by possible worlds—where **S** forms a belief that **p**, this belief is true. In order for this condition to be met, it is clearly essential that skeptical possible worlds—i.e., possible worlds in which the skeptical hypotheses are true—are modally far off, since otherwise there would be a near-by possible world in which **S** forms a false belief about **p**, such as in the worlds in which **S** is a (handless) BIV.

(2) Given that the skeptical possible worlds are indeed modally far off, however, then it ought to be clear that **S**'s belief (about self) that 'oneself is not a BIV' will also be safe as well. After all, it will be true, and since there are no near-by possible worlds in which the skeptical hypotheses obtain, there will be no near-by possible world in which **S** forms this belief and **S**'s belief is false.

(3) Insofar as **S** has knowledge of everyday propositions, then there is no reason to think that **S** lacks knowledge of the denial of skeptical premises. Closure is retained.

At p. 120, Prichard turns to the anti-skeptical proposal of 'Neo-Mooreanism.' This thesis is that **S** can *know* the *denials* of skeptical hypotheses. In other words, the first premise of the argument is false. Moore thought it was 'common sense' that we know the denial of skeptical hypotheses; and we know these denials on the basis of our knowledge of everyday propositions. But Moore's argument ultimately failed to engage in what the philosophical problem really is. Prichard says (paraphrased, pp. 120-21):

What is common between 'Mooreanism' and 'neo-Mooreanism' is that both hold that **S** can know the denials of skeptical hypotheses. That is where their similarity ends. The neo-Moorean maintains that it is essential to have an explanation for how such knowledge is possible. The explanation that a neo-Moorean standardly provides is externalist, although there are internalist variants.

According to externalism, although it is true that **S** lacks good reflectively accessible grounds for believing not-BIV, this fact alone *should not* decide whether **S** has knowledge of ordinary propositions. The neo-Moorean typically argues that **S**'s ordinary beliefs just satisfy certain externalist epistemic desiderata.<sup>26</sup>

---

<sup>26</sup> This 'epistemic desiderata' apparently includes that the truth of **p** is the result of reliable belief-forming processes.

Consider whether S's belief in the denial of a skeptical hypothesis is safe, i.e., whether it is true not only in the actual world but also in all (or nearly all) nearby possible worlds as well.

The first thing to note on this score is that if we do have the widespread knowledge that we credit to ourselves—such that the world is pretty much as we take it to be—then this would almost certainly entail that skeptical possible worlds are far-off worlds. But if that is the case, then our beliefs in the denials of skeptical hypotheses are safe by default, since clearly there cannot be a near-by possible world where one continues to believe in the target position and yet that belief be false.

Prichard recites two paragraphs about his intuitions about possible worlds, the sensitivity principle, belief-forming methods,<sup>27</sup> the safety principle, and epistemic luck.

The second element of the neo-Moorean position is an explanation for why we are so tempted to think that we lack knowledge of the denials of skeptical hypotheses. The externalist neo-Moorean will argue that it is epistemically internalist intuitions that are motivating that knowledge of this kind is impossible. With an externalist epistemology, S's inability to offer good reflectively accessible grounds in favor of **p** is entirely compatible with S's knowing that **p**.

Prichard concludes that there is no easy response to radical skeptical problem. Existing responses, despite their manifest problems, are capable of further development.

In chapter 7 entitled 'Understanding and the Value of Knowledge,' Prichard says that it is widely held that 'knowledge is of distinctive value.' He posits the Value problem: Understanding why knowledge is distinctively valuable. Some have questioned whether it is valuable at all. It is a challenge to explain why knowledge is more valuable than true belief. If we are unable to account even why knowledge is more valuable than mere true belief, then the very project of answering the value problem is a lost cause.<sup>28</sup>

---

<sup>27</sup> What exactly is a 'belief-forming method'? This is a problem with regard to reliabilism. If this concept cannot be specified (in one way or another), it shouldn't be central to a theory.

<sup>28</sup> From the perspective here, the project of answering the value problem *is* a lost cause. Value assertions are prescriptive. 'Is x valuable?' 'Is knowledge valuable?' These are not true/false questions.

## "The Gettier Problem and Epistemic Luck" (2019)<sup>29</sup>

In this essay, Prichard reiterates some previous arguments (2007, 2009) and says that "philosophers often mean different things when they talk about the Gettier problem."

**(1) The analytic Gettier problem:** The problem of offering a plausible theory of knowledge (such that can *deal with* Gettier cases). Assumes knowledge exists and describes the necessary and sufficient conditions for its attainment.

**(2) The luck Gettier problem:** What is the anti-luck condition on knowledge? A core intuition is that knowledge cannot be lucky, in the specific sense of **veritic luck**: Given how *S formed a belief*, it shouldn't be a matter of luck that the belief is true.<sup>30</sup>

Prichard says (paraphrased, p. 97):

The initial discussions of Gettier problem focused on what kind of condition we should add to justified true belief to ensure that *S* is not in a Gettier situation, so as to have knowledge. It is effectively the search for the anti-Gettier condition on knowledge.<sup>31</sup>

What we are after is not a gerrymandered anti-Gettier condition, but rather to determine just what it is about the JTB account that ensures that it is unable to eliminate the kind of veritic epistemic luck at issue in Gettier-style cases. In turn this means understanding the *nature* of our anti-luck intuition about knowledge.<sup>32</sup> This project is broadening our focus on knowledge, and the beginning of a methodology of an anti-luck epistemology.

What other kinds of case trade on the anti-luck intuition? He examines the lottery problem. Even if the odds are that ticket *x* is a loser, *S* still doesn't know that it is a loser. But if *S* reads a reliable newspaper, *S* can know that the ticket is a loser, even if the odds that there has been a newspaper error is higher than the lottery odds. One upshot of the lottery problem is that knowledge is not straightforwardly a function of the probabilistic strength of *S*'s evidence.

---

<sup>29</sup> Prichard, Duncan 2019. 'The Gettier Problem and Epistemic Luck' in Stephen Hetherington (ed.) *The Gettier Problem*. New York: Cambridge University Press.

<sup>30</sup> 'Belief forming methods' continue to be discussed. This contrasts with assessing *S*'s reasoning.

<sup>31</sup> To repeat: The PE definition isn't seeking an anti-Gettier condition, in order for *S* to know *p*. The definition is seeking the conditions for when *S* knows and doesn't know *p*. We shouldn't seek a gerrymandered anti-Gettier condition.

<sup>32</sup> Conditions PE 3 & PE 4b fulfill Prichard's 'anti-luck' world-view intuition.

Prichard poses the lottery problem in terms of luck and counter-possibilities (p. 98, paraphrased):

If **S** forms a true belief that a ticket is a loser (on astronomical odds that it is a loser), it still *seems* a matter of *luck* that the belief is true. After all, **S** could have very easily formed a false belief on this basis—all that it would have taken is for a few colored balls to fall in a *slightly different* configuration and **S**'s belief would have been false.

In contrast, the **S** who forms a true belief that a ticket is a loser based on the reliably produced newspaper, doesn't seem to have a belief that is true merely as a matter of *luck*, even though the odds of being right on this basis are much lower. But national newspapers take great care in providing accurate results, and despite the odds, it doesn't seem to be an easy possibility that **S** is wrong if **S** forms a belief on this basis.

Prichard presents his analysis as a broad project of rendering a general anti-luck condition on knowledge. What is at issue in the anti-luck condition? As discussed in 2005, 2007, and 2012, he says that we need to proceed in three stages:

- (1) Develop an account of luck.
- (2) Develop the specific sense in which knowledge is incompatible with luck.
- (3) Put these parts together to determine the anti-luck condition on knowledge.

Prichard says that luck *seems to be* a modal problem. Roughly a modal account of luck:

**A modal account of luck:** For an event to be 'lucky' means that it is an event that (keeping relevant initial conditions fixed) could *very easily* have not occurred.<sup>33</sup>

This explains why lottery wins are lucky, in that there are plenty of very close possible worlds where **S** doesn't win. It also explains why it isn't lucky that the sun rose this morning, in that there is no close possible world where the sun doesn't rise (although both events are out of our control). Combine the 'modal account' of luck with notion of 'veritic luck' (given how **S** forms the belief, it is not a matter of luck that **S** forms a true belief), and we get:

**Anti-luck condition:** Given how **S** forms a belief, **S**'s belief should not be false in close possible worlds.

---

<sup>33</sup> A stipulative technical fixed definiens definition (3c) for a group resemblance concept (i.e., 'luck').



Based upon the standard 'similarity ordering' as constructed by Stalnaker (1968) and Lewis (1973), low probability events can happen in close possible worlds. S's winning the lottery is a modally close but low-probability event. Although the odds are against winning, *not much needs to change* about the actual world to make it happen (just a few colored balls need to fall down in a different configuration), and so the *possible world* where S wins the lottery *is close* to the *actual world*. This is why S (and community) don't place bets on events with similar astronomical odds (modally far-fetched events) such as someone swimming the Atlantic Ocean. For any S, the national lottery's slogan, 'It could be you' is intended with a modal reading rather than probabilistic. In the modal sense, it really *could* be you, because *very little would need to change* otherwise. If S construed the slogan in the probabilistic sense, this seems problematic, as realistically it *couldn't* be you, as the odds are astronomically against S (14 million to 1).<sup>34</sup>

Anti-luck epistemology explains what is going on in "the lottery problem." Knowledge is not a straightforward function of S's probabilistic evidence because S needs a true belief that is not subject to veritic luck. This is unpacked in terms of a modal theory of luck:

**Modal Theory of Luck:** S's *basis for belief* should not result in false belief in close possible worlds.

That's what goes wrong in a lottery case. S's forming a belief that ticket is a loser by just reflecting on the odds is a *veritically lucky* way of forming a belief since it could so very easily have led to a false belief. (S's ticket could have won, and S would still *have believed on this basis* that the ticket had lost). In contrast, reading the lottery result in a national newspaper is *not veritically lucky*, in that S *couldn't* easily have ended up with a false belief *via this process*. There are rigorous precise constraints which ensure that newspapers provide correct results (i.e., misprints are nearly impossible).

---

<sup>34</sup> The PE definition has an alternative explanation for the slogans 'it could be you that wins the lottery' and 'you never know that you won't win the lottery.' Despite the high probability that it is true that the ticket will lose, S's evidence for believing p isn't connected to any particular ticket. With a fair drawing S can never be in a position to possess relevant premises for knowing the ticket is a loser. Knowledge condition 3 is not satisfied. Even after S considers the statistical evidence and discards the small chance that the ticket may win, and S has a personally justified belief, S cannot *cannot* have a justified (truth-connecting) belief that the ticket will lose.

Anti-luck epistemology has provided an account of the anti-luck condition:

**Anti-luck condition:** Given how **S** forms a belief, **S**'s belief should not be false in close possible worlds.

On the modal account of luck, we can *order* the degree of luckiness of the target event in terms of how modally close the non-obtaining of this event is. Similar to 2009 (p. 39):

For example, in the actual world **S** is not shot by the sniper's bullet, which whizzes past nearby. But **S** is luckier to be alive in the world where it whizzes by a few inches from his head than in the world where it flies a few meters away. Luck thus admits of degrees, and so the anti-luck condition should admit of degrees. (p. 101).

We are intolerant of high levels of luck when it comes to knowledge, and tolerant of low levels of luck. In between there is a sliding scale of tolerance.

For example, we are opposed to knowledge in the lottery case (**S** forms the belief that the ticket lost, by reflecting on the astronomical odds) because we recognize that **S** could have very easily formed a false belief, unlike **S**'s reading of the newspaper lottery results (where **S** couldn't have easily formed a false belief). In between there is a sliding scale, leading to penumbral cases, where the level of luck involved is such that we aren't confident either way. (p. 101).

In a footnote, Prichard states that he now uses the term *epistemic risk* rather than *epistemic luck* (and hence now argues for an *anti-risk epistemology*). After some rumination about the safety/sensitivity debate, the safety condition is introduced:

**Safety Condition:** When **S** knows, **S**'s true belief must be such that it couldn't have easily been false (when formed on the same basis). (p. 103).

Gettier-style cases essentially involve unsafe beliefs. Given how **S** was forming a belief about the sheep in the field, **S** could very easily on this basis have formed a false belief (i.e., if the sheep that was hidden from view had wandered out of the field).

The same goes for **S**'s forming a belief that the ticket lost just by reflecting on the odds involved. There is a very close world where this basis for belief leads to a false belief (i.e., the world where a few colored balls fall in a slightly different configuration).

In contrast, S's forming a belief by reading results in a reliable newspaper will result in a safe belief. It is not an easy possibility that this basis for this belief will result in a false belief.<sup>35</sup>

Prichard argues extensively for the preferability of safety over sensitivity and says that 'safety' by its nature only considers what is going on in close possible worlds, one's *modal neighborhood* if you will (p. 104). Anti-luck epistemology has led to an account of the anti-luck condition such that knowledge demands true belief that could not have easily been false, essentially safety. The 'anti-luck condition' on knowledge has been determined and has thereby solved the 'luck Gettier problem.' Can this be used to solve the 'analytical Gettier problem' too?

Prichard says that it can't be the case that we understand knowledge as mere true belief that satisfies the anti-luck condition, a position he has termed *robust anti-luck epistemology*, as opposed to a *modest anti-luck epistemology*, that merely contends that the anti-luck condition is a necessary condition on knowledge (2009, chapters 3 & 4). He says that the reason for this is that there is another fundamental intuition about knowledge:

**The ability intuition:** When S knows it must be down to S's cognitive agency in some significant way that the belief is true. (p. 105).

He says (quoted):

It is, after all, *you* who acquired the knowledge, so it must be the case that you did something cognitively significant in forming the target belief. The ability intuition can easily seem like just the other side of the coin to the anti-luck condition. If your true belief isn't due to cognitive agency then, presumably, it must be down to luck, right? And if your belief is due to luck, then surely that must be because it's down to your cognitive agency? (p. 105).

He says (paraphrased):

The relationships between anti-luck and ability intuitions are complex. What is of particular interest is that there can be non-lucky cognitive successes that are not significantly attributable to S's cognitive agency. The anti-luck condition only ensures that there is a certain match between S's belief and the truth across the relevant close possible worlds, but S's belief can have this modal profile and yet this has nothing to do at all with S's exercise of cognitive agency. (p. 105).

---

<sup>35</sup> Can we not contemplate the possibility of a malicious newspaper computer hacker?

He provides a 'Helpful Demon' example (similar to an example in 2009, p. 49):

S is forming beliefs about a certain domain in an entirely unreliable fashion. S forms a belief about the current room temperature by randomly picking two single digit numbers and combining them into a double-digit temperature. But with a helpful demon, the actual temperature changes with S's belief to match that double-digit belief. Accordingly, S not only forms true beliefs using his method; this is a *guaranteed method* to form true beliefs.<sup>36</sup> (p. 105).

His analysis of the case study (paraphrased):

S's beliefs satisfy anti-luck condition (i.e., there is nothing at all lucky about the fact that the beliefs are true). S doesn't just get to the truth in the actual world, but also in all close possible worlds too, since whenever S forms a belief about the temperature the demon will ensure that this belief is true. Crucially, however S's true belief has nothing to do with S's cognitive agency, but rather has everything to do with the demon helper. S's belief satisfies the anti-luck intuition without thereby satisfying the ability intuition.<sup>37</sup> (p. 106).

He continues (paraphrased):

An immediate upshot of this point is that robust anti-luck epistemology is not tenable as a response to the analytical Gettier problem. S can satisfy the anti-luck condition in such a way that it doesn't satisfy the ability intuition. It follows that adding the anti-luck condition to the JTB condition will deal with Gettier-style cases but will not present us with a complete theory of knowledge. (p. 106).

We need a theory of knowledge that demands S's non-lucky success is significantly attributable to S's cognitive agency. This is what is missing in the above demon case. While S in this scenario is forming safe beliefs about the temperature, S's cognitive agency is playing no explanatory role in accounting for why S's beliefs are safe (the explanation falling on the interventions of the demon helper. (p. 106)

In conclusion, Prichard advocates an anti-luck virtue epistemology, where the "virtue" captures the contribution of S's cognitive agency. This position, he says, is not a conjunction of anti-luck condition and ability condition upon knowledge. A proposal of this kind could be Gettierized. What is sought instead is a conception of knowledge such

---

<sup>36</sup> This is a *stipulated method* for S to (necessarily) form true beliefs. The 'Helpful Demon' example is entirely fictional with no analogous kinds of examples in the real world.

<sup>37</sup> S has a true belief but conditions 3 and 4b are violated. Also, condition 4a is seemingly violated (in the real world). S doesn't have knowledge.

that S's non-lucky cognitive success is significantly attributable to S's manifestation of cognitive agency. There is an important explanatory connection between S's non-lucky cognitive success and S's manifestation of cognitive agency, such that the non-lucky success cannot be merely incidentally correlated with cognitive agency. Ideally this proposal would help us respond to the analytic Gettier problem, and to recognize that it does so by drawing on how one should respond to the luck Gettier problem.

### **“Varieties of Epistemic Risk” (2022)<sup>38</sup>**

The motivation for a shift to ‘anti-risk’ epistemology is explained. The reason for this shift is to better respond to some subtle details of a modal account of knowledge:

In recent work, I have argued that we should shift our attentions away from anti-luck epistemology and towards a slightly different, though overlapping proposal—viz., anti-risk epistemology (and thus, in terms of the theory of knowledge in particular, *anti-risk virtue epistemology*). The concepts of risk and luck are closely related, in that luck and risk ascriptions tend to go hand-in-hand. For example, one was lucky to have survived that plane crash. Similarly, one’s life was at a great deal of risk during the plane crash. As this example illustrates, a key difference between the two notions is perspective, in that luck tends to involve a backwards looking assessment of the event in question (i.e., once tragedy is averted, at least in the case of good luck), while risk tends to involve a forwards-looking assessment of that event. Crucially however, the modal profile of the target event is roughly the same. Juggling dynamite is risky because there is a close possible world where undertaking that activity will lead to the relevant risk event (being blown up). In contrast, juggling oranges is not risky at all, as there is no close possible world where this activity leads to the relevant risk event (and certainly not the risk event of being blown up). We thus get a *modal account of risk*, which closely parallels the modal account of luck.

### **Conclusion**

Prichard’s philosophical methodology and worldview intuitions about knowledge include the terms ‘luck,’ ‘safety,’ ‘ability,’ ‘virtue,’ and ‘possible worlds,’ and are a marked contrast to the PE analysis of knowledge. The PE definitional approach seeks a social scientific conceptual analysis, while Prichard’s approach is predominately metaphysical. It is left to the reader’s judgement as to which epistemic theory is best.

---

<sup>38</sup> Prichard, Duncan 2022. ‘Varieties of Epistemic Risk’ *Acta Analytica* 37(1): 9-23.