

Epistemic Luck

Epistemic luck is a generic notion used to describe any of a number of ways in which it can be accidental, coincidental, or fortuitous that a person has a true belief. For example, one can form a true belief as a result of a lucky guess, as when one believes through guesswork that “C” is the right answer to a multiple-choice question and one’s belief just happens to be correct. One can form a true belief via wishful thinking; for example, an optimist’s belief that it will not rain may luckily turn out to be correct, despite forecasts for heavy rain all day. One can reason from false premises to a belief that coincidentally happens to be true. One can accidentally arrive at a true belief through invalid or fallacious reasoning. And one can fortuitously arrive at a true belief from testimony that was intended to mislead but unwittingly reported the truth. In all of these cases, it is just a matter of luck that the person has a true belief.

Until the twenty-first century, there was nearly universal agreement among epistemologists that epistemic luck is incompatible with knowledge. Call this view “the incompatibility thesis.” In light of the incompatibility thesis, epistemic luck presents epistemologists with three distinct but related challenges. The first is that of providing an accurate analysis of knowledge (in terms of individually necessary and jointly sufficient conditions for “*S* knows that *p*,” where ‘*S*’ represents the knower and ‘*p*’ represents the proposition known). An adequate analysis of knowledge must succeed in specifying conditions that rule out all instances of knowledge-destroying epistemic luck. The second challenge is to resolve the skeptical paradox that the ubiquity of epistemic luck generates: As will become clear in section 2c, epistemic luck is an all-pervasive phenomenon. Coupling this fact with the incompatibility thesis entails that we have no propositional knowledge. The non-skeptical epistemologist must somehow reconcile the strong intuition that epistemic luck is not compatible with knowledge with the equally evident observation that it must be. The third challenge concerns the special skeptical threat that epistemic luck seems to pose for more reflective forms of knowledge, such as knowing that one knows. Each of these challenges will be explored in the present article.

Table of Contents

1. Epistemic Luck and the Analysis of Knowledge
 - a. The Incompatibility Thesis
 - b. The Justified-True-Belief Analysis of Knowledge
 - c. The Gettier Problem
 - d. Purported Solutions to the Gettier Problem
 - i. No False Grounds
 - ii. No Essential False Grounds
 - iii. Defeasibility Approaches
 - iv. The Externalist Turn
 - v. The Causal Theory of Knowing
 - e. Controversial Cases
 - i. Beneficial Falsehoods
 - ii. Misleading Evidence One Does Not Possess
 - iii. Impact of These Cases
 - f. Where Things Stand
2. The Paradox of Epistemic Luck
 - a. The Knowledge Thesis
 - b. The Incompatibility Thesis (Again)
 - c. The Ubiquity Thesis
 - d. The Skeptical Challenge
 - e. Rejecting the Incompatibility Thesis
 - f. Knowledge-Destroying Epistemic Luck
 - i. Evidential vs. Veritic Luck
 - ii. Justification-Oriented Luck
 - iii. Modal Veritic Luck
 - g. Second-Wave Anti-Luck Epistemologies
 - i. Sensitivity
 - ii. Safety
 - h. Paradox Lost
3. Epistemic Luck and Knowing that One Knows
 - a. Internalism, Epistemic Luck, and the Problem of Knowing that One Knows
 - b. Epistemic Luck and Reflective Knowledge
4. Conclusion
5. References and Further Reading

1. Epistemic Luck and the Analysis of Knowledge

There is no settled agreement as to how best to characterize the accidentality or fortuitousness of

an epistemically lucky true belief. Some have attempted to cash out the accidentality of epistemically lucky beliefs modally. For example, Mark Heller (1999) contends that person *S*'s belief that *p* is epistemically lucky (and hence not knowledge) if *p* is true in the actual world, but there is at least one world, in a contextually-determined set of possible worlds, where *S*'s belief that *p* is false. On Duncan Pritchard's modal characterization (2005), *S*'s belief is epistemically lucky if it is true in the actual world, but false in a majority of nearby possible worlds where *S* forms the belief in the same way. Others (Riggs 2007; Coffman 2007) insist that epistemic luck be cashed out in terms of a *lack of control* condition. Each of these proposals has been criticized in the literature. Despite the lack of agreement concerning the exact nature of epistemic luck, there is widespread agreement that epistemic luck is incompatible with knowledge.

One of the earliest recorded illustrations of knowledge-destroying luck can be found in Plato's *Theaetetus*. In this dialogue, Socrates inquires as to what knowledge is. When Theaetetus suggests that knowledge is *true belief*, Socrates quickly convinces him that he is mistaken by noting that a jury may luckily arrive at a true belief either as a result of the rhetorical skill of a jurist intent on getting a certain verdict or on the basis of unsubstantiated hearsay, and in either case, the lucky true belief would fall short of knowledge. The Socratic challenge posed in the *Theaetetus* is to specify what must be added to true belief to get knowledge. To meet that challenge, one must provide an analysis of knowledge that correctly identifies the conditions that are individually necessary and jointly sufficient for *S* to know that *p* (where '*S*' represents the knower and '*p*' represents the proposition known). As will become readily apparent in what follows, the possibility of epistemic luck makes the already difficult task of meeting the Socratic challenge all the more difficult.

a. The Incompatibility Thesis

Epistemologists have long agreed with Plato that epistemic luck is incompatible with knowledge. To see just how widespread commitment to the incompatibility thesis is, consider the remarks of just a few representative epistemologists. In *The Problems of Philosophy* (1912, p. 131), Bertrand Russell asks the question: "Can we ever *know* anything at all, or do we merely sometimes by good luck believe what is true?"—the implication being that lucky true belief is not knowledge. In *Theory of Knowledge* (1990, p. 12), Keith Lehrer stresses that knowledge requires more than lucky true acceptance: "If I accept something without evidence or justification . . . and, as luck would have it, this turns out to be right, I fall short of *knowing* that what I have accepted is true." In *Reasons and Knowledge* (1981, p. 31), Marshall Swain maintains that: "lucky guesses do not constitute factual knowledge." In his *Stanford Encyclopedia of Philosophy* article on the analysis of knowledge (2006), Matthias Steup expressly endorses the incompatibility thesis: "Let us refer to a belief's turning out to be true because of mere luck as *epistemic luck*. It is uncontroversial that knowledge

is incompatible with epistemic luck.” In his *Internet Encyclopedia of Philosophy* article on epistemology (2007), David Truncellito concurs: “a lucky guess cannot constitute knowledge. Similarly, misinformation and faulty reasoning do not seem like a recipe for knowledge, even if they happen to lead to a true belief.” In *Epistemic Luck* (2005, p. 1), Duncan Pritchard calls attention to “the seemingly universal intuition that *knowledge excludes luck*, or, to put it another way, that the *epistemic luck* that sometimes enables one to have true beliefs . . . is incompatible with knowledge.”

The nearly universal intuition that epistemic luck is incompatible with knowledge is rooted in compelling examples like the following:

Jack of Hearts

Dylan is an avid euchre player. One night between hands, the dealer asks Dylan which card he believes to be on the top of the freshly shuffled euchre deck. Dylan thinks for a moment and, recalling his fondness of bowers, comes to believe that the top card is the jack of hearts. After Dylan reports his belief, the dealer turns over the top card, which just so happens to be the jack of hearts.

Since the probability of the jack of hearts being the top card of a randomly shuffled euchre deck is $1/32$, it is just a matter of luck that Dylan’s belief was true. Dylan certainly didn’t *know* that the jack of hearts was the top card. He just happened to guess correctly, and knowledge requires more than lucky guesswork.

b. The Justified-True-Belief Analysis of Knowledge

Examples like Jack of Hearts clearly show that true belief is not sufficient for knowledge. What, then, must be added to true belief in order to get knowledge? Prior to 1963, most epistemologists maintained that justification is what is required to convert true belief to knowledge, and as a result, they endorsed the *justified-true-belief analysis of knowledge*:

(JTB) For any subject *S* and any proposition *p*, *S* knows that *p* if and only if:

(i) *p* is true,

(ii) *S* believes that *p* [*Bp*], and

(iii) *S* is justified in believing that *p* [*Jp*].

Fallibilists and infallibilists disagree about the kind of justification required by (iii). Infallibilists maintain that knowledge requires infallible justification. *Infallible justification* entails the truth of the proposition for which it is justification. Fallibilists, on the other hand, endorse a weaker justification requirement. They contend that the kind of justification requisite for knowledge need only render probable, but need not entail, that for which it is justification.

At first blush, it might look as if infallible justification holds the key to eliminating epistemic luck and is, thus, the kind of justification needed for knowledge. After all, if *S* believes that *p* on the basis of infallible truth-entailing justification for *p*, it is impossible for *S* to be mistaken with respect to *p*. Unfortunately, the legacy of infallibilism is nearly wholesale skepticism. The point can be demonstrated as follows: Since our evidence for the non-*cogito* contingent empirical propositions we believe never entails the truth of those propositions, it follows that if the kind of justification required for knowledge is infallible truth-entailing justification, then we are never justified in believing, and hence never know, that such propositions are true. An infallibilist justification requirement would go a long way toward eliminating epistemic luck, but it would do so at the cost of making empirical knowledge impossible—hardly an adequate non-skeptical solution to the problem of epistemic luck.

Recognizing the skeptical implications of infallibilism, most contemporary epistemologists have embraced fallibilism so that empirical knowledge remains at least in principle possible. Fallibilistic justification is thought to rule out epistemic luck by making one's belief extremely probable. When one's belief that *p* is extremely probable, it is not just a matter of luck that one's belief is true. Recall Jack of Hearts. Prior to the dealer's turning over the top card, Dylan has no evidence as to what the top card is. As such, it is extremely improbable that the top card is the jack of hearts. Consider how Dylan's epistemic situation changes after the dealer turns over the top card, and Dylan sees the jack of hearts. Now Dylan has good perceptual evidence that the card is the jack of hearts. Given his newly-acquired perceptual evidence, it is now extremely probable that the card is the jack of hearts, and as a result, it is no longer just a matter of luck that his belief that it is the jack of hearts is true. Granted, it is possible that a Cartesian evil demon could have caused Dylan to hallucinate the jack of hearts right as the dealer flipped over some other card (which illustrates that Dylan's perceptual evidence doesn't entail that the card is the jack of hearts), and so, his evidence doesn't eliminate all chance of error; but it does make the chance of error extremely low, and when error is extremely improbable, it is not simply a matter of luck that one's belief is true.

c. The Gettier Problem

Although the role of the justification condition in the JTB-analysis is to rule out lucky guesses as instances of knowledge, it remains possible, given any fallibilistic account of justification, to have a

justified belief that is only luckily true, a fact that went largely unnoticed until the publication of Edmund Gettier's seminal article "Is Justified True Belief Knowledge?" (1963). Therein, Gettier provides two compelling counterexamples to the JTB-analysis of knowledge. He dubs these examples "Case I" and "Case II." Both cases involve a person who justifiably comes to believe a true proposition by validly deducing it from a justified-but-false belief. Consider first Gettier's Case II.

Case II: Smith has good evidence for believing that Jones owns a Ford [*J*]. Indeed, Smith's evidence for thinking that Jones owns a Ford is at least as strong as the evidence that we typically have for thinking that our friends and family members own the cars they do. Smith's evidence consists of the following: As far back as Smith can remember, Jones has always owned a Ford; just that morning, Jones gave Smith a ride while driving a Ford; and Smith was with Jones a few months back when Jones purchased a Ford exactly like the one she was driving when she offered Smith the ride earlier that morning. Based on this evidence, Smith justifiably believes that Jones owns a Ford [*J*]. On the basis of her justified belief that *J*, Smith justifiably deduces and comes to believe the disjunction that either Jones owns a Ford or Brown is in Barcelona [*J or B*], despite having no idea of Brown's whereabouts. As it turns out, Jones no longer owns a Ford. She recently sold her Ford and is driving a rental. However, purely by coincidence, Brown happens to be in Barcelona. Obviously, it is just a matter of luck that Smith's justified belief that *J or B* is true. Nearly every epistemologist who has considered this case agrees that Smith's *luck-infused* justified-true-belief that *J or B* falls short of knowledge.

Here is a slightly modified version Gettier's other example. **Case I:** While waiting for a job interview, Smith sees Nelson take the coins out of her pocket, count them (ten coins in all), and then put them back in her pocket. Smith also overhears the boss on the phone telling someone that Nelson is the person who will get the job. On the basis of this evidence, Smith justifiably believes the conjunction:

(*N*) Nelson will get the job, and Nelson has ten coins in her pocket.

On basis of her justified belief that *N*, Smith deduces and justifiably comes to believe:

(*P*) The person who will get the job has ten coins in her pocket.

Despite Smith's evidence, *N* is false. The boss misspoke on the phone. Actually, it is Smith, not Nelson, who will get the job, and purely by chance, Smith happens to have exactly ten coins in her pocket. Once again, it is just a matter of luck that Smith's belief—this time her belief that *P*—is true. With these two examples, Gettier showed that fallibilistic justification is incapable of eliminating all forms of knowledge-destroying epistemic luck and that, as a result, justified true belief is not

sufficient for knowledge.

d. Purported Solutions to the Gettier Problem

Gettier's paper gave rise to a plethora of articles attempting to solve the problem that now bears his name. Many of these purported solutions sought to resolve the problem by supplementing the JTB-analysis with a fourth condition, while others abandoned the JTB-analysis in favor of non-traditional justificational accounts of knowledge. Consider first some of the prominent fourth condition responses.

i. No False Grounds

In both of Gettier's examples, Smith justifiably infers a true belief from a justified-but-*false* belief, and it has seemed to many that a true belief is not knowledge when it is deduced from a false belief. As a result, a number of epistemologists sought to resolve the Gettier problem by supplementing JTB with a "No False Grounds" clause along the following lines:

(NFG) *S* knows that *p* if and only if (i) *p* is true, (ii) *S* believes that *p*, (iii) *S* is justified in believing that *p*, and (iv) *S*'s justification for *p* does not rest on any false beliefs.

An analysis of knowledge can be too strong or too weak: It is *too strong* if it is possible for a person to know that *p* without satisfying all of the conditions spelled out in the analysis. It is *too weak* if one can fail to know that *p* when all the conditions in the analysis have been met. To see that NFG is too strong, we need only modify Gettier's Case II as follows:

Café

Smith is sitting in a café in Barcelona *with Brown* having a cup of espresso. While there, with Brown, Smith justifiably infers and comes to believe that *J or B* on the basis of her justified-but-false belief that Jones owns a Ford [*J*] and on the basis of her justified-true-belief that Brown is Barcelona [*B*].

In this scenario, Smith has excellent evidence for *B* along with her misleading evidence for *J*. Since Smith knows *B* is true and validly deduces *J or B* from her knowledge that *B*, it is not at all lucky that her justified belief that *J or B* is true; and so, Smith knows that *J or B*, despite the fact that part of her evidence, namely, *J*, is false. Hence, NFG is false, for it entails that a person fails to know that *p* whenever *any* part (even a dispensable and thus superfluous part) of her justification is false, when, intuitively, a person with some false evidence for *p* can still know that *p* provided she

has at least one independent chain of all-true-evidence justification for p .

ii. No Essential False Grounds

In Gettier's Case II (where Smith clearly fails to know that J or B), Smith's justification for J or B essentially depends on Smith's justified-but-false belief that J . In Café (where intuitively Smith does know that J or B), Smith has two independent strands of justification for J or B . The first strand is her justified-but-false belief that J . The second strand is her justified-true-belief that B . As a result, Smith could dispense with the first strand entirely and still be justified in believing the disjunction J or B . Our markedly different appraisals of Smith's epistemic status vis-à-vis J or B in these two cases suggest that the presence of false grounds precludes the knowledge that p only when those grounds play an indispensable role in a person's justification for p . Given this insight, it might seem that the *no false grounds* condition in NFG should be replaced with a *no essential false grounds* condition as follows:

(NEFG) S knows that p if and only if (i) p is true, (ii) S believes that p , (iii) S is justified in believing that p , and (iv) S 's justification for p does not *essentially* depend on any false beliefs.

Unfortunately, NEFG is too weak because there can be *all-true-evidence* Gettier cases—cases where the person's justification for her lucky true belief does not depend on any false beliefs. An example of an all-true-evidence Gettier case is provided by Brian Skyrms's (1967) case involving Sure-Fire matches:

Pyromaniac

Pete knows that Sure-Fire matches have always lit in the past when struck. Pete also knows that the match he is holding is a Sure-Fire match. Based on this evidence, which he knows to be true, Pete justifiably believes that the match he is holding will light when struck [L]. Unbeknownst to Pete, the match he is holding is a defective Sure-Fire match (the first ever!) with impurities that raise its combustion temperature above that which can be produced by striking friction. As luck would have it, just as Pete strikes the match, a sudden burst of Q-radiation ignites the match.

In Pyromaniac, Pete has a justified true belief that L , which is based entirely on *true* evidence that Pete *knows*, and yet, it is still just a matter of luck that his belief is true. This example shows that one can have a lucky true belief that p that falls short of knowledge, even when all of one's evidence for p is true. Thus, NEFG is too weak. There are genuine Gettier cases that it fails to rule out.

iii. Defeasibility Approaches

In each of Gettier's original cases, there is a true proposition unbeknownst to Smith such that were that proposition added to the rest of Smith's evidence, Smith would no longer be justified in believing the Gettiered belief. Call such a proposition a *defeater*. In Case I, the defeater is the true proposition that Nelson will not get the job [$\sim N$]. If $\sim N$ were added to Smith's evidence, Smith would not be justified in believing that the person who will get the job has ten coins in her pocket, for she would no longer have any idea who will get the job. In Case II, the defeater is the true proposition that Jones does not own a Ford [$\sim J$]. Since Smith has no knowledge as to Brown's whereabouts, if $\sim J$ were added to Smith's evidence, she would no longer be justified in believing that J or B . Notice, however, that in the case of Café (where Smith is with Brown in Barcelona), the true proposition $\sim J$ is not a defeater, because adding $\sim J$ to Smith's evidence in Café would not prevent Smith from being justified in believing that J or B . Smith would still be justified in believing J or B on the basis of her justified true belief that B .

Defeasibility theorists contend that a person fails to know that p whenever there is a defeater for her justification for p . Their proposal for solving the Gettier problem is to supplement the JTB-analysis with a *No Defeaters* condition as follows:

(ND) S knows that p if and only if (i) p is true, (ii) S believes that p , (iii) S is justified in believing that p , and (iv) there are no defeaters for S 's justification for p .

The biggest problem facing the No Defeaters approach is that there is no agreement among defeasibility theorists themselves as to the correct account of defeaters. For example, Roderick Chisholm (1964) and Peter Klein (1971) have characterized defeaters as follows:

(D1) When evidence e justifies S in believing that p , then proposition d is a *defeater* for S 's justification if and only if (i) d is true and (ii) the conjunction of d and e does not justify S in believing that p .

Keith Lehrer and Thomas Paxson Jr. (1969) contend that D1 is too weak, as a definition of defeaters, because it counts as defeaters certain statements that intuitively are not defeaters. They offer the following case in point:

Grabit

While at the library, I see a student of mine, Tom Grabit, take a book from the shelf, conceal it under his coat, and leave the library without checking it out. I know Tom Grabit well, and I am

sure that he stole the book. I justifiably believe that Tom Grabit stole the book, and he did.

Intuitively, I know that Tom Grabit stole the book. But here's the rub: Unbeknownst to me, Tom Grabit's mother said that on the day in question Tom was not in the library, indeed, he was thousands of miles away, and that Tom's identical twin, John Grabit, was in the library. On the D1 account of defeaters, the following true proposition is a defeater for my justification for thinking that Tom stole the book:

(*M*) Tom's mother said that Tom was not in the library at the time of the theft, but his identical twin John was.

If *M* were added to my evidence, I would no longer be justified in believing that Tom stole the book. This result might seem like the right result until we discover that Tom's mother is both delusional and a pathological liar, that she said these things to herself in her padded cell, that John Grabit is a figment of her demented mind, and that Tom stole the book just as I thought. Lehrer and Paxson argue that the fact that it is true that a delusional mental patient uttered false statements about Tom Grabit's location on the day of the theft should not defeat my knowledge that Tom stole the book. They conclude that the Chisholm/Klein account of defeaters should be replaced with the following account:

(D2) When evidence *e* justifies *S* in believing that *p*, then proposition *d* is a *defeater* for *S*'s justification if and only if (i) *d* is true, (ii) *S* is completely justified in believing that *d* is false, and (iii) the conjunction of *d* and *e* does not justify *S* in believing that *p*.

In Grabit, I do not have any evidence concerning what Tom's mother said or didn't say, and so, I am not completely justified in believing that it is false that she said those things. As a result, condition (ii) of D2 is not satisfied, and so, on the Lehrer/Paxson account of defeaters that fact that Ms. Grabit said those things is *not* a defeater for my evidence that Tom stole the book. Consequently, on the D2-account of defeaters, I have an undefeated justified true belief that Tom stole the book and thus know that Tom stole the book, which is the intuitively correct result.

The Chisholm/Klein D1-account of defeaters gets the Grabit case wrong, for it entails that the true statement *M* defeats my justification for thinking that Tom stole the book. Since *M* is a defeater on the Chisholm/Klein account, their account entails that I do not know that Tom stole the book; when, intuitively, I do know that Tom stole the book. I saw him steal it, and the insane ramblings of his demented mother do nothing to undermine my knowledge.

Now consider a different case:

Locked

John Lock is compulsive when it comes to locking his doors. This morning when he left for work, he locked the front door and tripled checked that the door was locked. It is now 11:00 a.m., and John is sitting in his office recalling his morning ritual. He clearly and distinctly remembers locking his front door and triple checking to make sure that it was in fact locked. On the basis of this vivid memorial evidence e , he comes to believe that his front door is locked. Lucy Lock, John's wife, is notoriously unreliable about locking the doors when she leaves home, which is why John always insists on leaving home after Lucy leaves for work. Unbeknownst to John, Lucy forgot her workout clothes and returned home at 10:30 a.m. to retrieve them, and she just happened to lock the door when she left five minutes later.

So, at 11:00 a.m., John's belief that the front door is locked is true. Presumably, John does not know that the front door is locked. He thinks the door is locked because he remembers locking it, but that is not why it is locked. It is locked because Lucy absentmindedly and uncharacteristically happened to lock it on her way out. Intuitively, John's knowledge is defeated by the following true proposition:

(U) The door was unlocked by Lucy at 10:30 a.m.

If U were added to John's memorial evidence e , John would no longer be justified in believing that his front door is locked. On the Chisholm/Klein D1-account of defeaters, U is a defeater because U is true and the conjunction of U and e would not justify John in believing that his front door is locked. However, on the Lehrer/Paxson D2-account of defeaters, U would not count as a defeater because sitting in his office at 11:00 a.m., John has no evidence concerning whether or not his wife returned home to retrieve her gym clothes, and so, he is not completely justified in believing it false that the door was unlocked by Lucy Lock at 10:30 a.m. Since U is not a defeater on the Lehrer/Paxson account, their account entails that John knows that his front door is locked; when, intuitively, he fails to know that his door is locked, because it is just a matter of luck that Lucy absentmindedly locked it when she left.

The problem for the No Defeaters approach, then, is this: D1 is too weak of an account of defeaters, and as a result, employing a D1-account of defeaters in ND would make ND too strong an account of knowledge; whereas D2 is too strong an account of defeaters, and so, employing it in ND would make ND too weak. Absent an adequate account of defeaters, the No Defeaters approach fails to provide a solution to the problem of epistemic luck.

iv. The Externalist Turn

Externalist theories of justification maintain that epistemic justification is (at least) partly a function of features *external* to the cognizer, that is, features outside the cognizer's ken. For example, one prominent externalist theory, process reliabilism, makes a belief's justificatory status a function of the *actual* reliability (rather than the *perceived* reliability) of the process producing that belief. One motivation behind externalism with respect to justification is its unique ability to provide a truth connection that conceptually links justification with truth. To appreciate the importance of this motivation, recall that the role of the justification condition in the JTB-analysis is to rule out lucky guesses as instances of knowledge. In order for justification to be able to properly play that role, there must be some sort of internal connection between justification and truth that makes the former objectively indicative of the latter. Indeed, many epistemologists insist that it is precisely its internal connection to truth that distinguishes *epistemic* justification from *moral* and *pragmatic* justification. To be objectively indicative of truth, justification must be conceptually connected (not merely coincidentally or contingently connected) with truth. In order for there to be a conceptual connection between justification and truth, the following condition must hold: In every possible world W , if conditions C make belief b justified in W , then conditions C also make b objectively probable in W . The rationale for requiring such a truth connection is this: If there were no conceptual connection between justification and truth, it would be just as much a matter of luck when a justified belief turned out to be true as when an unjustified belief turned out to be true. Moreover, a better justified belief would be no more likely to be true than a much less well justified belief, for without a truth connection, no amount of justification is an objective indication of truth.

Unlike externalist theories of justification, no internalist theory of justification can provide the desired conceptual connection between justification and truth. Internalist theories maintain that epistemic justification is solely a function of states *internal* to the cognizer, such as perceptual states, belief states, memorial seemings, and introspective states. Examples of such theories include classical foundationalism, coherentism, and evidentialism. That no internalist theory of justification can provide a truth connection can be demonstrated as follows: Every internalist theory of justification maintains that the conditions that make a belief justified are entirely specifiable in terms of states internal to the cognizer, 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 's belief b internalistically justified in W_D do not make b objectively probable in W_D , no internalist theory is capable of providing a truth connection. Because internalistic justification is not conceptually connected to truth, one can easily be internalistically justified in holding a false

belief, which can in turn be used to justifiably infer some other belief that may coincidentally turn out to be true. Consequently, employing an internalistic justification condition in the JTB-analysis makes JTB particularly susceptible to Gettier cases.

At first glance, externalistic justification looks more promising as a means of preventing luck from playing a role in the acquisition of true belief, for some externalist theories of justification do provide a conceptual connection between justification and truth. Consider, for example, the following simplified version of process reliabilism:

(PR) S 's belief b is justified in world W if and only if S 's belief b is produced by a belief-forming cognitive process [BCP] that is W -reliable (where a W -reliable BCP is a process that tends to produce beliefs in W that are true in W).

Since PR asserts that a belief is justified in W if and only if it is produced by a W -reliable BCP, and since, by definition, the beliefs produced by a W -reliable process tend to be true in W , PR-justified beliefs have a high objective probability of being true. Because reliably-produced, externalistically-justified beliefs are objectively likely to be true, one might think that replacing the internalistic justification condition in the traditional JTB-analysis with an externalistic justification condition would render JTB immune to Gettier-style counterexamples. William Harper (1996) quickly dispels any such notion, with the following counterexample:

Falcon

Smith believes that Jones owns a Ford. Smith was with Jones when Jones purchased her Pinto; Smith has seen the official title to the car in Jones's name; Jones is a reliable informant that has never deceived anyone; and just this morning, Jones gave Smith a ride to work in her Pinto. Smith has a reliably-produced and reliably-sustained belief that Jones owns a Ford. It is now 1:00 p.m. Unbeknownst to Smith, at noon Jones's Pinto was vaporized by a terrorist bomb; but, also unbeknownst to Smith, exactly at noon Jones won a Ford Falcon in the State Lottery. Hence, Smith has a reliably-formed true belief that Jones owns a Ford, but her belief is not knowledge.

While internalistic justification may be particularly susceptible to being undermined by Gettier-style knowledge-destroying luck, Harper's counterexample shows that the Gettier problem plagues all fallibilistic theories of justification, both internalistic and externalistic alike. Whatever virtues externalistic justification might have, solving the Gettier problem is not among them.

v. The Causal Theory of Knowing

In his early work, Alvin Goldman (1967) offers a different diagnosis of what has gone wrong in Gettier cases. In Case II, what makes *J or B* true is the fact that Brown is in Barcelona, but this fact plays no causal role in Smith's coming to believe that *J or B*. In Case I, what makes *P* (*P* = the person who will get the job has ten coins in her pocket) true is the number of coins in Smith's pocket, but this fact plays no role in Smith's coming to believe *P*. What causes Smith to believe *P* is the fact that Nelson has ten coins in her pocket, and this latter fact is not what makes *P* true. Goldman observes that in these cases there is no *causal connection* between the Gettiered belief and the fact that makes it true. It is the absence of such a connection that allows for the possibility of belief's being true purely by luck. Goldman concludes that the traditional JTB-analysis should be replaced with the following *causal theory of knowledge*:

(CTOK) *S* knows that *p* if and only if the fact that *p* is causally connected in an appropriate way with *S*'s believing that *p*.

The appropriate knowledge-producing causal processes that Goldman identifies include: (i) perception, (ii) memory, (iii) inferentially reconstructed causal chains, each inference of which is warranted, and (iv) combinations of (i)-(iii). The causal theory correctly handles all of the cases we have considered so far. We have already seen how it handles Gettier's original cases. In Café, what makes *J or B* true is the fact that Brown is in Barcelona, and that fact is appropriately causally connected with Smith's believing that *J or B*, because Smith is having an espresso with Brown in Barcelona. Accordingly, CTOK correctly entails that, in Café, Smith knows that *J or B*. In Grabit, I see Tom steal the book. Tom's stealing the book in plain eyesight perceptually causes me to believe that he did, and so, once again, CTOK yields the right result: I know that Tom Grabit stole the book. In Locked, what makes it true that the front door is locked is the fact that Lucy locked it, and this fact plays no causal role in John's believing his front door is locked. In this case, CTOK correctly entails that John does not know that his front door is locked. He's just lucky that Lucy happened to lock it. Finally, in Falcon, what makes it true that Jones owns a Ford is the fact that she just won a Ford Falcon in the state lottery, and that fact plays no causal role in Smith's believing that Jones owns a Ford, and so, Smith fails to know that Jones owns a Ford.

Despite its success in handling these cases, the causal theory falls prey to the following counterexample:

Fake Barn County

An eccentric farmer in Minnesota owns all of the land in Fake Barn County. Wanting to appear

much richer than he is, this farmer has erected fake barns all throughout the county. From the road, these fake barns look exactly like real barns, when, in reality, they are just two dimensional barn façades. While nearly every barn-looking structure in the county is a fake, there are a few real barns interspersed among the myriad fakes. Henry, who is driving through Fake Barn County, has no idea that there are any fake barns in the county. Looking out the window of his car, Henry sees what looks to be a barn on the hill just up the road and comes to believe that there is a barn on the hill. Purely by chance, Henry happens to be looking at one of the few real barns in the county.

Intuitively, Henry does not know that there is a barn on the hill. He is just lucky to be looking at one of the few real barns in the county. The lucky nature of his present belief becomes even more obvious once we discover that Henry has been forming barn beliefs ever since entering Fake Barn County, and all of these other barn beliefs have been false. Henry has consistently been duped by the façades.

The causal theory fails because it cannot account for Henry's lack of knowledge in this case. Henry is now looking at one of the few real barns in the county, and this real barn is what is *causing* him to believe that there is a barn on the hill. Since Henry's true belief that there is a barn on that hill is appropriately caused via perception by that very barn on that hill, the causal theory mistakenly entails that Henry knows there is a barn on that hill, when clearly he does not.

e. Controversial Cases

As analyses of knowledge aimed as at solving the Gettier problem have grown in sophistication and complexity, so have the purported counterexamples aimed at refuting these analyses. Some of these purported counterexamples are sufficiently complex and controversial that there is no consensus among epistemologists as to whether or not the person in the example knows the proposition in question. Two such cases are discussed below.

i. Beneficial Falsehoods

Although the *no essential false grounds* approach was largely abandoned once it was shown that there can be *all-true-evidence* Gettier cases—cases where *S*'s justification for her lucky belief *p* does not depend on any false beliefs—there has remained nearly universal agreement among epistemologists that a person fails to know that *p* if her justification for *p* *essentially* depends on a false belief. Peter Klein (2008) is a noteworthy exception. He contends that there can be beneficial falsehoods—falsehoods essential to one's justification that nevertheless give one knowledge. Here is one example that Klein offers in support of his controversial view:

Appointment

Based on memory, I believe that my secretary told me on Friday that I have an appointment with a student on Monday. Based on that justified-but-false memorial belief, I come to justifiedly believe that I have an appointment on Monday. As it turns out, I do have that appointment on Monday, and my secretary did tell me of the appointment. However, he didn't tell me on Friday. He told me on Thursday.

Klein contends that I know that I have an appointment on Monday [*A*], even though my belief that *A* essentially depends on my false belief that my secretary told me *on Friday* about my Monday appointment. It might seem that the false belief that my secretary told me on Friday that I have a Monday appointment is not essential to my justification for *A*, because if I “remember” that my secretary told me *on Friday* of my Monday appointment, then presumably I also actually remember that my secretary *told me* I have an appointment on Monday, and this latter belief is true. But suppose that the secretary was out with the flu the first three days of the week, and also suppose that I do not remember being on campus on Thursday. In fact, I'm confident that I wasn't on campus on Thursday, having totally forgotten that I briefly stopped in on Thursday to get my mail. Klein contends that in such a situation I would not believe that my secretary told me of the appointment at all, unless I believed that he told me this *on Friday*. Klein contends that I know that *A*, even though that belief essentially depends on my false belief that my secretary told me *on Friday* about my Monday appointment.

What distinguishes beneficial falsehoods from knowledge-destroying falsehoods? Under what circumstances does a false belief *f* allow *S* to acquire knowledge that *p*? Klein's answers to these questions are rooted in and flow out of his preferred theory of knowledge. Klein contends that knowledge requires both propositional and doxastic justification. Proposition *p* is *propositionally justified* for *S* if and only if *S* has an *epistemically adequate basis* for *p*. *S*'s belief that *p* is *doxastically justified* for *S* if and only if *S*'s belief that *p* has an *appropriate causal basis*. The basic idea is that in order for *S* to know that *p*, *S*'s belief that *p* must be epistemically justified and appropriately caused. Armed with the distinction between propositional and doxastic justification, Klein argues that a false belief *f* is a *beneficial falsehood* just in case the following seven conditions are met: (i) *f* is false, (ii) *S*'s belief that *f* is doxastically justified (that is, *S*'s belief that *f* has an appropriate causal pedigree), (iii) the belief that *f* is essential in the causal production of the belief that *p*, (iv) *f* propositionally justifies *p*, (v) *f* entails a true proposition *t*, (vi) *t* propositionally justifies *p*, and (vii) whatever doxastically justifies *S* in believing that *f* also propositionally justifies *t* for *S*. When these conditions are met, Klein contends that the false belief *f* is “close enough to the truth” to give one knowledge that *p*. Applied to the case at hand, *F* is the false proposition that my secretary told me *on Friday* that I have an appointment on Monday, and *T* is the true proposition

that my secretary *told* me that I have an appointment on Monday. Klein contends that my belief that *F* meets all the conditions for being a beneficial falsehood: (i) *F* is false, (ii) my belief that *F* is doxastically justified (appropriately caused) by the fact that he did tell me, (iii) my belief that *F* is essential to my believing that *A* (for if I didn't believe *F*, I would not believe he had told me about an appointment at all and so I would not believe *A*), (iv) *F* propositionally justifies *A*, (v) *F* entails the true proposition *T*, (vi) *T* propositionally justifies *A*, and (vii) the fact that my secretary told me on Thursday about my Monday appointment propositionally justifies *T* for me.

Klein contends that, in the case at hand, it doesn't really matter what day my secretary told me that I have an appointment on Monday. What matters is the fact that he told me. The false belief *that he told me on Friday* is close enough to the true proposition *that he told me* as to give me knowledge that I have an appointment on Monday.

While interesting and provocative, Klein's case is difficult to assess because it depends on controversial assumptions about belief individuation. Is it possible, for example, to believe that my secretary told me on Friday that I have an appointment on Monday [*F*], without also believing that my secretary told me that I have an appointment on Monday [*T*]? If not, then rather than providing us with a case of an indispensable knowledge-generating false belief, Klein may have simply given us another case of justificatory over-determination; for if it is impossible to believe *F* without also believing *T*, then there seem to be two independent strands of justification only one of which depends on a false belief, in which case Appointment is simply an analogue of Café above.

ii. Misleading Evidence One Does Not Possess

Gilbert Harman (1973) contends that *S*'s knowledge that *p* can be undermined by readily available misleading evidence that *S* does not possess. In Harman cases, despite the fact that the undermining evidence is misleading, if *S* were to possess it, *S* would no longer be justified in believing that *p*. The idea behind Harman cases seems to be this: Since the misleading evidence is *readily available*, it is just a matter of luck that *S* does not possess that evidence, and since luck is incompatible with knowledge, *S* fails to know that *p*. Here is one of Harman's cases:

Assassination

A political leader is assassinated. His associates, fearing a coup, decide to pretend that the bullet hit someone else. On Nationwide television they announce that an assassination attempt has failed to kill the leader but has killed a secret service man by mistake. However, before the announcement is made, an enterprising reporter on the scene telephones the real story to his newspaper, which has included the story in its final edition. Jill buys a copy of that paper,

reads the story of the assassination, and believes that the President has been assassinated based on the story. What she reads is true, and so are her assumptions about how the story came to be in the paper. (1973, 143)

Harman insists that Jill does not know that the President has been assassinated. He finds it highly implausible that Jill should know simply because she lacks evidence that everyone else possesses. Harman's diagnosis is that Jill's knowledge is undermined by readily available evidence – the misleading televised retraction – that she does not possess.

Epistemologists who have reflected on Harman's Assassination case remain divided over whether or not Jill knows that the President has been assassinated. Those who think that she does know that the President has been assassinated tend to focus on the facts that (i) all of her evidence is true, (ii) she knows her evidence is true, and (iii) the evidence she has is an accurate indicator of the President's assassination.

Those epistemologists who think that Jill does not know that the President has been assassinated do not focus on the *quality* of Jill's evidence, which is impeccable. Rather, they focus on the *lucky nature* of her evidence. If Jill had turned on the TV when she got home, like she usually does, she would have seen the televised retraction, and she would have found herself in the same epistemic predicament as everyone else. Given the conflicting reports, she would not have known what to believe. Clearly, Jill is lucky to be in the evidential situation she is in. Since luck is generally thought to be incompatible with knowledge, these epistemologists conclude that Jill fails to know that the President has been assassinated.

iii. Impact of These Cases

Controversial cases like these make the challenge of providing an accurate analysis of knowledge even more difficult. If Jill does know that the President has been assassinated, then all those theories of knowledge that imply that she lacks such knowledge (including Harman's own theory) are mistaken. On the other hand, if Jill does not know that the President has been assassinated, then all those theories that imply she does are mistaken. Similarly, if I do not know that I have an appointment on Monday, then all those theories that imply I do (including Klein's theory) are mistaken. If I do know that I have an appointment on Monday, then all those theories that imply I lack such knowledge are mistaken. The competing intuitions these cases engender make the already difficult task of arriving at a mutually agreed upon account of knowledge even more formidable.

f. Where Things Stand

While various proponents of the above proposals might still embrace them, the general consensus is that none of the above attempts at eliminating epistemic luck succeeds. One problem with these first attempts at resolving the Gettier problem is that they tended to emerge in a piecemeal fashion as responses to specific counterexamples, only to fall prey to more elaborate counterexamples themselves. What seems to be needed is a better understanding of epistemic luck itself. If we can get clear on the exact nature of knowledge-destroying luck, we might be in a better position to formulate a condition that can eliminate it. The next section will examine a number of attempts at clarifying the nature of knowledge-destroying luck and will assess several modal conditions that have been proposed to eliminate such luck.

2. The Paradox of Epistemic Luck

In addition to generating problems for those epistemologists seeking an analysis of knowledge, the phenomenon of epistemic luck gives rise to an epistemological paradox in its own right. The paradox is generated by the following three theses: the knowledge thesis, the incompatibility thesis, and the ubiquity thesis. The paradox arises because each of these theses is antecedently plausible, but together they form an inconsistent triad. Each thesis is discussed below.

a. The Knowledge Thesis

According to *the knowledge thesis*, we possess a great deal of knowledge about the world around us. Commonsense tells us that the knowledge thesis is true. For example, I know that I am in a coffee shop. I know that I am drinking a cup of coffee. I know that I am wearing a blue shirt. I know that I am typing on a laptop computer. And I know that the person sitting next to me is talking on his cell phone at an inappropriate volume. You know that you have eyes. You know that you are reading an IEP article on epistemic luck. You know that the article you are reading is written in English. Together, we know a lot. At least, we think we do, until we encounter a skeptical paradox like the paradox of epistemic luck.

b. The Incompatibility Thesis (Again)

The incompatibility thesis is the thesis that epistemic luck *simpliciter* is incompatible with knowledge. As noted above, there has been nearly universal agreement among epistemologists that knowledge is incompatible with epistemic luck. The post-Gettier literature is replete with evermore-sophisticated counterexamples to the array of purported accounts of knowledge proffered in an effort to resolve the Gettier problem. The standard formula for generating a

counterexample to a purported analysis of knowledge is to conjure up a case where, despite satisfying all the conditions in the analysis, it is still just a matter of luck that the person's belief that p is true. The element of luck involved is *ipso facto* thought to prevent the belief from being an instance of knowledge. The nearly unanimous acceptance of such examples illustrates just how widespread commitment to the incompatibility thesis is.

c. The Ubiquity Thesis

Epistemic luck is an all-pervasive phenomenon that infects every fallibilistic epistemology in one form or other. Its inescapability can be demonstrated as follows: To convert true belief to knowledge, every viable fallibilistic 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 evil 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 is always in some sense a matter of luck when a *merely* internalistically justified belief turns out to be true. To see why, consider my twin in an evil demon world W_D . By hypothesis, he has the same beliefs that I have, he has the same memorial seemings that I have, he possesses the same experiential evidence that I possess, and he goes through exactly the same internal reflections that I do. In short, our internal cognitive lives are phenomenologically, doxastically, and reflectively indistinguishable. Consequently, if I satisfy the internal conditions for justifiedness (whatever they may be), then my demon-world twin satisfies them as well, and so, we are both internalistically justified in our beliefs. If, on the other hand, I fail to satisfy those conditions, then my twin also fails to satisfy them, and so neither of us is internalistically justified in our beliefs. Now assume the former scenario where both of us are internalistically justified in our beliefs. The only relevant difference between my twin and me is that he is being systematically deceived, whereas, as epistemic good fortune would have it, I am not. If he and I were to change places, there would be no introspectable difference, and each of us would continue to believe as we do, only now I would be the hapless victim of demon deception. Clearly, I am epistemically fortunate to be in the world I am in (assuming I am in the world I take myself to be in) and not in W_D . Since I am lucky to be in the world I am in, there is a clear sense in which it is epistemically lucky that my internalistically justified beliefs are true. My twin is not nearly so lucky, for, thanks to the demon, all of his internalistically justified beliefs are false. Since these results can be generated no matter which internalistic theory of justification one employs, it is always a matter of luck when a merely internalistically justified belief happens to be true.

Truth-connected externalist approaches (for example, reliabilist, truth-tracking, and safety-based

accounts) avoid this kind of epistemic luck. However, they are subject to another kind of ineliminable epistemic luck. Recall, from section 1, the externalistic process-realist account of epistemic justification:

(PR) S 's belief b is justified in world W if and only if S 's belief b is produced by a belief-forming cognitive process that is W -reliable.

Call a belief that is justified in virtue of being reliably produced a *PR-justified belief*. Although it is not typically a matter of luck when a PR-justified belief turns out to be true (since PR-justification is conceptually connected to truth), it is a matter of luck when a belief turns out to be PR-justified. To see why, consider, once again, my twin in the demon world W_D . By hypothesis, he and I share the same beliefs, possess the same evidence, go through the same internal reflections, and have phenomenologically, doxastically, and reflectively indistinguishable cognitive lives. Even so, our beliefs do not have the same PR-justificatory status. His beliefs are not PR-justified, because they are produced by processes that the demon has rendered unreliable in W_D , whereas my beliefs are PR-justified because they are produced by processes that are reliable in the actual world (Again, I'm assuming, for the sake of the example, that the actual world is the world we think it is.). According to PR, it is not a matter of luck that my beliefs are true and his beliefs are false, because my beliefs are PR-justified and his are not, and PR-justified beliefs have a high objective probability of being true. What is a matter of luck is the fact that my beliefs are PR-justified and his are not. After all, we both take ourselves to be in non-demon-manipulated worlds, and we both take ourselves to have reliably-produced PR-justified beliefs. As luck and ill luck, respectively, would have it, I am correct and he is incorrect. Since there is no introspectively discernible difference between our worlds, given what each of us has to go on, there is a clear sense in which I could have just as easily been mistaken and been the one with demon-rendered-unreliable processes. Compared to my twin, I am epistemically fortunate to be in a non-demon world where my cognitive faculties are reliable. Since I am epistemically lucky (compared with my twin) to be in a world where I have reliable cognitive processes, there is clearly a sense in which it is just a matter of luck that I have PR-justified beliefs.

Analogous considerations can be applied to any externalistic constraint on knowledge. Consider the externalistic condition of being a *safe* belief (to be explained below). While a safe belief's being true is not epistemically lucky, having safe beliefs is epistemically lucky, for in a demon world none of one's beliefs are safe. Since every fallibilistic epistemology incorporates either an internalistic justification condition or an externalistic condition, no fallibilistic epistemology can rid us of epistemic luck's intractable presence.

d. 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, 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 observation that it must be. Since the ubiquity thesis is unassailable, the anti-skeptical epistemologist must reject the incompatibility thesis.

e. 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 is true, 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 but 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.

f. Knowledge-Destroying Epistemic Luck

Unger has successfully identified three types of harmless epistemic luck, but not all forms of epistemic luck are benign. What is needed is an account of knowledge-undermining luck.

i. Evidential vs. Veritic Luck

Mylan Engel Jr. (1992) distinguishes two kinds of epistemic luck, evidential luck and veritic luck, and argues that only the latter is incompatible with knowledge. Engel characterizes these two types of luck as follows:

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

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

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. 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.

Engel argues that all genuine Gettier cases involve veritic luck. In Gettier's Case II presented above, Smith's belief that J 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 is just a matter of luck that Smith's belief that J 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 the truth of p in a way that makes it entirely coincidental from S 's point of view that p is true.

Engel then uses the distinction between evidential and veritic luck to assess Harman cases. Jill is not veritically lucky in believing that the President has been assassinated, for she has accurate, reliable evidence concerning the assassination in the form of a reputable newspaper's column, and given this evidence, it is not a matter of luck that her belief is true. However, Jill is evidentially lucky—she is lucky to be in the evidential situation that she is in, for had she turned on the TV and seen the fabricated retraction, she would have been in a much worse evidential situation vis-à-vis the President's assassination. Lucky for her, she did not turn on the TV. Like the bank teller above, Jill is lucky to have the evidence she does, but given her evidence, she is not lucky that her belief is true. Having argued that only veritic luck is incompatible with knowledge, Engel concludes that Jill does know the President has been assassinated. If Engel is right, then Harman cases do not provide examples of knowledge-undermining luck.

ii. Justification-Oriented Luck

Hamid Vahid (2001) maintains that there are two types of knowledge-destroying epistemic luck. He agrees with Engel that veritic luck as characterized by VL is incompatible with knowledge, but he argues, *contra* Engel, that there is a kind of evidential luck (which he dubs 'justification-oriented luck') that is also incompatible with knowledge. Vahid contends that knowledge-precluding justification-oriented luck is a function of how easily a person's belief could have been unjustified:

(JL) A person suffers from knowledge-precluding *justification-oriented luck*, when she is justified in believing that p , but given her epistemic circumstances, she could have easily been

unjustified in holding that very belief.

Vahid contends that Harman's Assassination case provides an example of knowledge-precluding justification-oriented luck. Jill could have *easily* been unjustified in believing that the President was assassinated. Had she turned on the TV like she usually does, she would not have been justified in holding that belief. Vahid concludes that Jill does not know that the President was assassinated—her knowledge is destroyed by justification-oriented luck.

While JL might yield the right result in Harman's Assassination case, it seems to yield the wrong result with respect to the bank teller case. The teller is justified in believing that the bank president is the robber because she just happened to look up during the brief moment when his mask had slipped and clearly saw the robber's face, but she could have easily been unjustified in this belief. Had she continued to look in her cash drawer while nervously collecting the cash for the robber, she would not have seen the robber's face. Clearly, the teller knows that the bank president is the robber, and yet, JL implies that she lacks such knowledge.

iii. Modal Veritic Luck

Duncan Pritchard (2003) agrees that, of these types of luck, only veritic luck is incompatible with knowledge, but he replaces Engel's 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 if and only if S 's belief that p is true in the actual world a but false in nearly all nearby possible worlds in which S forms the belief in the same manner as in a .

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.

g. Second-Wave Anti-Luck Epistemologies

The post-Gettier literature is rife with attempts at supplementing or amending the traditional JTB-analysis with a satisfactory anti-luck constraint on knowledge. As surveyed in Section 1, the first wave of proposals included adding a no-false-grounds or no-*essential*-false-grounds condition to

JTB, supplementing JTB with a defeasibility condition, incorporating an externalistic justification condition in JTB, and replacing JTB with a causal theory of knowing. These and similar proposals have fallen prey to ever-more-complicated Gettier-style examples. The general consensus is that none of these proposals succeeds. Second-wave luck-eliminating proposals invoke counterfactual or subjunctive constraints on knowing, principal among them: *sensitivity* and *safety*. Let us consider each of these proposals in turn.

i. Sensitivity

S's belief that *p* is *sensitive* to *p*'s truth-value if and only if *S* would not believe that *p* if *p* were false (that is, if and only if *S* does not believe *p* in any of the closest $\sim p$ -worlds). To be sure, sensitive belief does preclude veritic luck, but it does so at a steep price. First, the sensitive-true-belief account of knowledge results in closure failure. Second, there are compelling reasons to think that knowledge does not require sensitivity. Let's examine each cost in turn.

Most epistemologists regard it as all but axiomatic that we can expand our knowledge by competently deducing some currently unknown proposition *u* from some other known proposition *k* whenever we know that *k* entails *u*. This widely-embraced idea is codified in the *principle of epistemic closure* which has been formulated in each of the following ways:

(PEC1) If *S* knows that *p* and *S* knows that *p* entails *q*, then *S* knows (or is in a position to know) that *q*.

(PEC2) If (i) *S* knows that *p*, (ii) *S* knows that *p* entails *q*, (iii) *S* competently deduces *q* from her knowledge that *p* and that *p* entails *q*, and (iv) *S* comes to believe *q* as a result of that competent deduction, then *S* knows that *q*.

One reason the principle of epistemic closure has enjoyed such widespread endorsement is this: If I know that *p* and know that *p* entails *q* and I deduce and come to believe *q* from that knowledge, my belief that *q* could not be false (because knowledge is factive and the truth of *p* entails *q*, together with the truth of *p*, guarantees the truth of *q*).

The following example illustrates why sensitive-true-belief accounts of knowledge result in closure failure. I currently believe that I am in a coffee shop [*C*]. My belief that *C* is sensitive. If I were not in a coffee shop, I would not believe that I was, for if I were not in a coffee shop, I would be somewhere else, for example, the grocery store or my office, and would not mistakenly think that I was at a coffee shop. Since my belief that *C* is sensitive (that is, I would not believe it if it were false), the sensitive-true-belief account of knowledge entails that I know that *C*. I also currently

believe that I am not at home in bed having a lifelike dream of being in a coffee shop [$\sim H$], but my belief that $\sim H$ is not sensitive, for if I were at home in bed having a lifelike dream of being in a coffee shop (that is, if $\sim H$ were false), I would still believe that $\sim H$. So, according to the sensitive-true-belief account, I do not know that $\sim H$. Of course, my being at the coffee shop entails that I am not at home in bed dreaming that I am in a coffee shop (that is, $C \implies \sim H$), and I know that $C \implies \sim H$. The sensitive-true-belief account results in closure failure because it entails that I know that C and know that $C \implies \sim H$, but I do not know (and cannot come to know) that $\sim H$ on that basis.

Most epistemologists regard the principle of epistemic closure to be so plausible that they find any theory of knowledge that results in closure failure deeply problematic if not outright absurd. In fairness to sensitivity theorists, they recognize that their theories entail closure failure and acknowledge the antecedent implausibility of closure failure, but they argue that, despite its counterintuitiveness, there are principled reasons for thinking that knowledge is not closed under known implication. They grant that we have all sorts of ordinary knowledge, but insist that we do not know and *cannot* know that the skeptic's hypotheses are false. Thus, they embrace closure failure because they think that it accurately captures our actual epistemic situation. Perhaps sensitivity theorists are right, but given how widely accepted the principle of epistemic closure is, it would be preferable to identify an anti-luck constraint that avoids closure failure.

The second major problem facing the sensitivity proposal, as Jonathan Vogel (1999) shows with *Hole-In-One*, is that knowledge does not require sensitivity. The fourth hole at Augusta National Golf Course where The Masters is played is a tricky 240-yard par 3, euphemistically called "Flowering Crabapple." In 2007, not one player shot a hole-in-one on this diabolical hole, and there were only eleven birdies throughout four rounds of play. Right now, I know that not all seventy-two players in this year's Masters will shoot a hole-in-one on Flowering Crabapple in the first round of play, but my belief to this effect is not sensitive. Were every golfer to shoot a hole-in-one on Flowering Crabapple in Round One of the Masters in defiance of the astronomical odds against it, I would still believe that they were not going to do so. So, sensitivity is not necessary for knowledge.

ii. Safety

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

Peter Murphy (2005) employs Saul Kripke's famous counterexample to sensitivity to show that strong safety results in closure failure. Suppose the following is true of Fake Barn 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 of the façades are red. Driving through Fake Barn County, Mary is unaware that the 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, for 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 Mary knows she's looking at a *blue barn*, but does not know she's looking at a *barn*.

Weak safety is open to a different worry. If knowledge only requires weakly-safe justified true belief, then a person who justifiably believes her lottery ticket will lose *knows* that her ticket will lose (unless, of course, it happens to win), because in the overwhelming majority of nearby worlds, her ticket is a loser. Many epistemologists (though not all) insist that people do not know their lottery tickets will lose, prior to hearing the announced results. Anyone convinced that people do not know their tickets will lose, before learning of the results, will think that weak safety is too weak of an anti-luck constraint on knowledge.

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. Next, have Smith justifiably infer a disjunction of the form $J \text{ or } \sim G$, where Smith has no evidence whatsoever that $\sim G$ is true and where unbeknownst to Smith, $\sim G$ is true in all nearby worlds. Let $\sim G =$ Brown will not win a Grammy. Suppose that, unbeknownst to Smith, Brown is totally devoid of musical talent and there is no remotely close world where Brown wins a Grammy. Then, Smith's true belief that $J \text{ or } \sim G$ will be safe, but veritically lucky nonetheless, because given Smith's evidence, it is just a matter of luck that $J \text{ or } \sim G$ is true. Since the safe-true-belief account cannot rule out all cases of veritic luck, safe-true-belief is not sufficient for knowledge.

Hiller and Neta's example also shows that Pritchard's modal account of veritic luck [MVL] is not the correct analysis of veritic luck. Smith's belief that $J \text{ or } \sim G$ is clearly veritically lucky: Smith bases her belief that $J \text{ or } \sim G$ on her justified-but-false-and-unsafe-belief that Jones owns a Ford. Since Smith has absolutely no knowledge or evidence of Brown's total lack of musical talent, given Smith's evidence, it is just a matter of luck that her belief that $J \text{ or } \sim G$ is true. But MVL entails that Smith's belief is not veritically lucky. According to MVL, a belief is veritically lucky if it is true in the actual world but false in nearly all nearby worlds where Smith forms the belief in the same manner. In Hiller and Neta's case, Smith's belief that $J \text{ or } \sim G$ is true in the actual world, but it is

also true in all nearby worlds where it is formed in the same way (because $\sim G$ is true in all nearby worlds). Thus, according to MVL, Smith's belief that J or $\sim G$ is not veritically lucky. Since Smith's belief is veritically lucky, the MVL analysis of veritic luck is mistaken.

h. Paradox Lost

The paradox of epistemic luck dissolves once we recognize that the incompatibility thesis is false. Not all forms of epistemic luck are incompatible with knowledge. Certainly propositional, existential, and facultative luck are compatible with knowledge, and at least some forms of evidential luck, like the evidential luck had by the bank teller above, are also compatible with knowledge. There is growing consensus that veritic luck is the principal form of knowledge-destroying luck. Since veritic luck is far from ubiquitous, the incompatibility of veritic luck with knowledge poses no general threat to the possibility of knowledge. One can know that p whenever it is not a matter of veritic luck that one's justified belief that p is true.

3. Epistemic Luck and Knowing that One Knows

Although there remains broad disagreement over how exactly to formulate the condition needed to rule out knowledge-destroying epistemic luck in a theory of knowledge, there is widespread consensus that whatever the correct condition is, S does not need to know that that condition has been met in order to know that p . The point can be illustrated as follows: Let the expression " S is not Gettiered with respect to p " serve as a placeholder for whatever the correct substantive luck-eliminating condition is. If we add this condition to the traditional justified-true-belief analysis, we get the following schema for analyzing knowledge:

(K) S knows that p if and only if:

(i) p is true,

(ii) S believes that p ,

(iii) S is justified in believing that p , and

(iv) S is not Gettiered with respect to p .

According to (K), S does not need to *know* that conditions (i)-(iv) are met in order to know that p . All that (K) requires for S to know that p is that conditions (i)-(iv) *be* met. Since S need not know or even believe that she is not Gettiered with respect to p in order to know that p , the possibility of

Gettier-style, knowledge-destroying, veritic luck poses no special obstacle to first-order knowledge (where ‘first-order knowledge’ refers to *knowing* that p and ‘second-order knowledge’ refers to *knowing that one knows* that p). As long as S is not veritically lucky with respect to p , she will know that p , according to schema (K), provided she has a justified true belief that p .

The situation seems to be quite different when it comes to *knowing that one knows*, for one of the most natural ways of coming to know that one knows that p is by knowing that one has met the conditions required for knowing that p , and knowing the latter requires knowing that one is not Gettiered with respect to p . The burden of the present section is to examine whether the phenomenon of knowledge-destroying epistemic luck undermines more reflective forms of knowledge, such as, knowing that one knows.

a. Internalism, Epistemic Luck, and the Problem of Knowing that One Knows

Arch internalist H.A. Pritchard (1950) famously remarked: “We must recognize that whenever we know something we either do, or at least can, by reflecting, directly know that we are knowing it.” Other internalists have been less sanguine about the prospects of second-order knowledge. For example, Roderick Chisholm (1986) argues that one cannot generally know that one knows on the grounds that one cannot generally know whether or not one’s evidence for p is defeated by Gettier considerations. Is Chisholm right? Does the Gettier problem pose special—indeed, generally insurmountable—obstacles to internalistically knowing that one internalistically knows that p ? Richard Feldman (1981) does not think so. He thinks that the Gettier problem poses a minor obstacle to second-order knowledge, but one that can be easily overcome with minimal intellectual effort. Mylan Engel Jr. (2000) disagrees. Siding with Chisholm, Engel argues that the Gettier problem poses three distinct challenges to second-order knowledge, which, when taken together, threaten to undermine the possibility of knowing that one knows. Michael Roth (1990) contends that the Gettier problem poses no threat to second-order knowledge whatsoever. To assess these competing views, it will be helpful to have a clearer idea of just what is required for internalistic knowledge.

Post-Gettier internalists with respect to knowledge tend to work within the JTB+ tradition in that they maintain that, in addition to true belief, knowledge requires internalistic justification as well as some fourth externalistic anti-luck condition to rule out Gettier cases. Accordingly, by replacing condition (iii) in schema (K) above with an explicitly internalistic justification condition, we arrive at a schema for internalistic knowledge that most internalists would readily embrace:

(K_i) S internalistically knows (knows_i) that p if and only if:

(k1) p is true,

(k2) S believes that p ,

(k3) S is internalistically justified (justified _{i}) in believing that p , and

(k4) S is not Gettiered with respect to p .

Since (K _{i}) provides a perfectly general account of knowledge _{i} , we can arrive at the conditions for second-order knowledge _{i} simply by substituting S knows _{i} that p for p in schema (K _{i}):

(K _{i} K _{i}) S knows _{i} that S knows _{i} that p if and only if:

(kk1) S knows _{i} that p is true,

(kk2) S believes that S knows _{i} that p ,

(kk3) S is justified _{i} in believing that S knows _{i} that p , and

(kk4) S is not Gettiered with respect to S knows _{i} that p .

Chisholm doubts that (kk3) can be satisfied. To appreciate Chisholm's worry, consider one of the most natural and straightforward ways of satisfying condition (kk3), namely, being justified _{i} in believing that one has met all of the conditions required for knowing _{i} that p :

(J _{i} K _{i} p) S is justified _{i} in believing that S knows _{i} that p if and only if:

(jk1) S is justified _{i} in believing that p ,

(jk2) S is justified _{i} in believing that S believes that p ,

(jk3) S is justified _{i} in believing that S is justified _{i} in believing that p , and

(jk4) S is justified _{i} in believing that S is not Gettiered with respect to p .

Since (jk1) is identical to (k3), (jk1) is satisfied whenever S knows _{i} that p ; and if we assume both doxastic and justificational _{i} transparency (that is, if we assume that we have introspective access to what we believe and to our justification _{i} for what we believe), as do many internalists, then (jk2)

and (jk3) also pose no special problems for the would-be second-order knower.

Chisholm's concern is with (jk4). He contends that we cannot typically tell whether or not our evidence for p is defeated by Gettier considerations. Based on this contention, Chisholm argues as follows: Let S be any one of us and let p be a proposition that S knows_i. Since S cannot tell whether S 's evidence for p is defeated by Gettier considerations, S is not justified_i in believing that S is not Gettiered with respect to p . Hence, (jk4) is not satisfied. Since (jk4) is not satisfied, S is not justified_i in believing that S knows_i that p , that is, (kk3) is not satisfied. Since (kk3) is a necessary condition for knowing_i that one knows_i that p , S does not know_i that S knows_i that p . The gist of Chisholm's argument is this: Since we are not justified_i in believing that we are not Gettiered with respect to p , we do not know_i that we know_i that p .

Feldman disagrees. He thinks that, with minimal effort, a person who knows_i that p can be justified_i in believing that she is not Gettiered with respect to p . Feldman offers two reasons for thinking that it is relatively easy to be justified_i in believing that one's evidence for p is not defective and thus that one is not Gettiered with respect to p . Since Feldman is primarily concerned with determining when a person who knows_i that p knows_i that she knows_i that p , he assumes that S has first-order knowledge_i that p when presenting his reasons.

No False Evidence

Assume that S knows_i that p . S 's knowing_i that p entails that S is justified_i in believing that p . Since S is justified_i in believing that p , S is also justified_i in believing that all of her evidence for p is true. Since false evidence is usually what makes one's evidence defective, S is justified_i in believing that her justification_i for p is not defective and thus that she is not Gettiered with respect to p .

Induction

Since S has very rarely found herself to be the victim of Gettier-type situations, she is justified_i in believing that such situations are very rare and atypical. Given their rarity and atypicality, S is justified_i in believing that she is not in such a situation with respect to p .

Feldman contends that No False Evidence and Induction provide S with good internalistic reasons for believing that she is not Gettiered with respect to p . Since internalistic justification is a function of having good internalistic reasons and S has good internalistic reasons for believing that she is not Gettiered with respect to p , S is justified_i in believing that she is not Gettiered with respect to p , that is, (jk4) is satisfied. Since (jk1)-(jk3) are also easily satisfiable, with minimal intellectual effort,

S can be justified _{i} in believing that she knows _{i} that p . Feldman concludes that satisfying (kk3) poses no special obstacle to knowing _{i} that one knows _{i} .

Engel contends that the Gettier problem generates three distinct challenges for the would-be second-order knower—challenges that threaten to undermine the satisfaction of (kk1), (kk3), and (kk4), respectively:

(1) *First-order actual Gettierization*: One way the Gettier problem can preclude S from knowing _{i} that she knows _{i} that p is by preventing S from knowing _{i} that p . If S is Gettiered with respect to p , then S fails to know _{i} that p , and thus, she fails to know _{i} that she knows _{i} that p , since (kk1) is unsatisfied. However, since first-order actual Gettierization precludes second-order knowledge _{i} only when it obtains, it poses no greater threat to second-order knowledge _{i} than it poses to first-order knowledge _{i} .

(2) *First-order possible (but non-actual) Gettierization*: While *actual* first-order Gettierization, when it obtains, undermines second-order knowledge _{i} by falsifying (kk1), *possible* (but non-actual) first-order Gettierization threatens to thwart one of the most natural ways of satisfying (kk3), namely, satisfying conditions (jk1)-(jk4) of $J_iK_i p$. Like Chisholm, Engel's concern here is with (jk4). He argues that the reasons Feldman offers—No False Evidence and Induction—do not provide adequate reasons for thinking that one is not Gettiered with respect to p . No False Evidence is not a good reason to think that one has not been Gettiered with respect to p because, as noted in Section 1, there can be *all-true-evidence* Gettier cases, a point that Feldman himself demonstrated in an earlier article (Feldman, 1974). While No False Evidence may provide S with a reason for thinking that she is not the victim of a Gettier case involving a justified-false-belief, it provides her with no reason to think that she is not the victim of an *all-true-evidence* Gettier case. The problem with Induction is that many of the Gettier cases described in the literature are what we might call “invisible” Gettier cases, that is, they are cases such that, were they to obtain, the Gettier victim would never find out. They are cases that look and feel like knowledge and pass away unnoticed. Unless Pyromaniac Pete is wearing a Geiger counter, he will never discover that it was Q-radiation and not striking friction that caused his defective Sure Fire match to light. Unless John Lock interrogates Lucy Lock about her morning routine, he will likely never discover that she unlocked the doors to their house at 10:30 a.m. Unless Henry leaves the highway and investigates, he will likely never discover that most of the barn-looking structures are façades. Considerations such as these make it plausible to think that invisible Gettier cases are more likely to be the norm than visible Gettier cases. The fact that S has rarely found herself to be Gettiered in the past may provide her with a reason for thinking that visible Gettier cases are rare, but it provides her no reason to think that invisible Gettier cases are rare, and without such a reason, she is not justified _{i} in believing that she is not being (invisibly) Gettiered with respect to p . Engel concludes that it is

much more difficult to be justified_i in believing that one is not Gettiered with respect to p than Feldman alleges.

(3) *Meta-Gettierization*: Engel dubs second-order Gettierization “meta-Gettierization.” Just as first-order Gettierization occurs when S ’s justification for p is defective in a way that makes S veritically lucky with respect to p , meta-Gettierization occurs when S ’s justification for believing that S knows_i that p is defective in a way that makes S veritically lucky with respect to S knows_i that p . By way of illustration, Engel asks us to consider Professor Cleaver, a fictitious philosophy professor from the 1950s, who, as a pre-Gettier epistemologist, justifiably accepts the JTB-analysis of knowledge. Since Cleaver is justified_i in believing that knowledge_i is justified_i true belief, he is justified_i in believing that (jk1)-(jk3) are jointly sufficient for being justified_i in believing that one knows_i that p . Since he is justified_i in believing that (jk1)-(jk3) are jointly sufficient for being justified_i in believing that one knows_i that p , he is justified_i in believing that he knows_i that p provided that he is justified_i in believing that he has justified_i-true-belief that p . Let p be a proposition that Cleaver knows_i. If he *believes* that he knows_i that p , and if he is *justified_i* in believing that he knows_i that p on the basis of his justified_i-but-*false*-belief that *knowledge_i is justified_i true belief*, together with his justified_i-true-belief that he has a justified_i-true-belief that p , then Cleaver will have a justified_i-true-belief that he knows_i that p , which falls short of knowledge_i because his justification_i essentially depends on his justified_i-but-*false*-belief that *knowledge_i is justified_i true belief*. The point generalizes. Anytime that Cleaver comes to believe that he knows_i that p on the basis of his justified_i-but-*false*-belief that the JTB-analysis is correct, he will automatically be meta-Gettierized and will, thus, fail to know_i that he knows_i that p . Engel then argues that whether those of us who have grown up in the post-Gettier enlightenment can avoid Cleaver’s fate depends on whether any of us justifiedly_i believes a *true* epistemology. Since no epistemology to date is immune to objection, Engel thinks it doubtful that any of us holds a true epistemology (no matter how well justified_i we might be in our preferred epistemology). Given how likely it is that we are operating with a false epistemology, Engel contends that whenever we come to believe that we know_i that p on the basis of our preferred epistemology, we are almost certain to become yet another meta-Gettier casualty, for we are almost certain to have based our belief that we know_i that p on a justified_i *false* belief about the requirements for knowledge_i.

Roth contends that the debate over whether the Gettier problem poses a major or minor obstacle to second-order knowledge_i is entirely misguided. He argues that Gettier considerations pose no obstacle to second-order knowledge_i whatsoever. His argument is rooted in what he calls the *Fallibilist Assumption Governing Empirical Knowledge*:

(FA) For every proposition of the form Kp (where p is empirical and K is the knowledge operator), there are certain contingencies such that: (i) their obtaining is physically possible,

(ii) were they to obtain, Kp would be false, and (iii) S is completely justified in disregarding any of these contingencies in considering whether she has adequate justification for p .

Roth contends that there are two types of Kp -falsifying contingencies. “Type I contingencies” satisfy conditions (i), (ii), and (iii) of (FA). “Type II contingencies” satisfy conditions (i) and (ii), but not (iii). Roth asks us to imagine a great dividing wall – The Wall of Fallibilism – that separates the Type I contingencies from the Type II contingencies. As Roth envisions it, the Wall of Fallibilism plays an important role in protecting us from knowledge-destroying epistemic luck. If, given S 's evidence for p in circumstances C , it is simply a matter of luck that p is true in C , then S does not know that p in C . To ensure that it is not just a matter of veritic luck that S 's belief that p is true (in C), S must be *suitably protected* from error with respect to p (in C). According to Roth, the Wall protects us from the slings and arrows of outrageous Type I error possibilities by cordoning us off from these remote *properly ignorable* Kp -falsifying contingencies. We do not need evidence that these contingencies do not obtain in order to know_i that p . Being safely outside the Wall, we do not need to take them into account in our epistemic reflections at all. Their sheer remoteness and improbability protects us from having to worry about them. As long as they do not actually obtain, these contingencies provide no obstacle to knowledge_i whatsoever. But the Wall does not provide us with all the protection from luck and error that we need in order to possess knowledge_i. We must also be protected from error with respect to those Type II contingencies that are inside the Wall. These p -falsifying contingencies are genuinely in doubt. Were any of these contingencies to obtain, p would be false, and as a result, so too would Kp . To protect us from these realistic non-ignorable $\sim p$ -possibilities, we need justification that precludes them. The picture of fallible knowledge_i that emerges is this:

S knows_i that p only if (i) S 's justification_i is strong enough to rule out all of the relevant Type II $\sim p$ -possibilities inside the Wall and (ii) none of the Type I contingencies outside the Wall obtain.

Roth thinks that the Wall metaphor explains why Gettier considerations pose no obstacle to second-order knowledge_i. Gettier considerations are paradigm cases of Type I contingencies. We do not need to know_i or even believe that Type I contingencies do not obtain in order to know_i that p . As long as no Type I contingencies obtain, S will know_i that p provided she satisfies the other conditions required for knowing_i that p . Like Type I contingencies generally, Gettier considerations only undermine knowledge_i when they obtain. We do not need to know_i or even believe that no Gettier circumstances obtain in order to know_i that p . As long as they do not obtain, we will know_i that p provided we have met the other conditions required for knowing_i that p . Since Gettier contingencies are outside the Wall, Roth contends that it is perfectly proper to ignore them when trying to determine whether one knows_i that p .

Roth's reason for thinking that Gettier contingencies pose no obstacle to knowing_i that one knows_i is that he thinks that Gettier possibilities are properly ignorable Type I contingencies that lie safely outside the Wall. The problem with Roth's argument is that the Wall's location is not fixed. As Roth himself admits, where the Wall is situated is relativized to a particular attempt to acquire knowledge_i of a particular proposition. Which contingencies are outside the Wall and which are not, that is, which contingencies are properly ignorable and which are not, is a function of the proposition one is attempting to come to know_i and the circumstances under which one is trying to come to know_i it. While Gettier contingencies vis-à-vis p are clearly properly ignorable where coming to know_i that p is concerned, they are not properly-ignorable when it comes to knowing_i that Kp . To the contrary, it seems that Kp -destroying Gettier contingencies are precisely the kind of contingencies that one needs to be able to rule out in order to know that *one knows that p* . Gettier contingencies are not p -falsifying contingencies (for p is true in Gettier situations), but they are Kp -falsifying contingencies. As such, they are Type I contingencies when it comes to knowing_i that p , but Type II contingencies when it comes to knowing_i that *one knows_i that p* . In effect, the Wall moves outward where second-order knowledge_i is concerned. The very same Gettier contingencies that are outside the p -Wall are inside the Kp -Wall. Being inside the Kp -Wall, they are not properly ignorable when it comes to knowing_i that Kp . To know_i that *one knows_i that p* , one must know_i that no Gettier Kp -falsifying contingencies obtain. It is precisely because we cannot generally know_i that no Gettier contingencies obtain that Chisholm and Engel contend that second-order knowledge_i is difficult to attain.

b. Epistemic Luck and Reflective Knowledge

Even if veritic luck poses no special problem for reflectively knowing that one knows, Duncan Pritchard contends that another more worrisome kind of epistemic luck does preclude such knowledge. *Reflective epistemic luck* arises when, *from the agent's reflective position*, it is just a matter of luck that her belief is true. More precisely:

MRL For all S and p , the truth of S 's belief that p is *reflectively lucky* if and only if S 's belief that p is true in the actual world but, in nearly all nearby possible worlds *consistent with what S is able to know by reflection alone*, were S to believe p , p would be false.

When it comes to modal reflective luck, the epistemically relevant possible worlds are ordered in a non-standard way solely in terms of what the agent is able to know on the basis of her subjective internal reflections alone. Accordingly, any possible world consistent with S 's having that same internally accessible evidence that she has in the actual world will be *reflectively* equally close to the actual world. Since, by hypothesis, S would have exactly the same internally accessible evidence in a demon world or a BIV-world that she has in the actual world, these worlds are just as close,

reflectively, to the actual world as is the world where everything is just as it seems. Since our ordinary commonsense perceptual beliefs are false in a wide variety of these reflectively equally close skeptical-scenario possible worlds, Pritchard maintains that MRL entails that our ordinary commonsense perceptual beliefs—if true in the actual world—are reflectively lucky. [Whether MRL actually entails that all of our true commonsense perceptual beliefs are reflectively lucky is by no means obvious. The fact that our commonsense beliefs are false in malevolent demon and BIV worlds does not show that these beliefs are false in *nearly all* reflectively equally close possible worlds. After all, for every malevolent demon world where we are systematically deceived, there is a corresponding benevolent demon world that is just as close, reflectively, in which the benevolent demon sees to it that all of our commonsense beliefs are true.]

Pritchard thinks that reflective luck is not incompatible with ordinary knowledge (he thinks only veritic luck is), but he insists that reflective luck is incompatible with a much-desired internalistic kind of robust *reflective knowledge*. Pritchard contends that skeptical challenges force us to confront the fundamental human epistemic predicament, to wit, that we cannot know, on the basis of reflection alone, that the skeptic's radical hypotheses are false. For example, he thinks that we cannot know, by reflection alone, that we are not bodiless brains being kept alive in vats of nutrient being deceived into thinking we have hands.

If Pritchard is right that we lack reflective knowledge that the skeptic's hypotheses are false, then those who think that reflective knowledge is closed under known entailment face an even greater skeptical threat. According to the *principle of epistemic closure* (PEC₁): If *S* knows that *p* and also knows that *p* entails *q*, then *S* either knows or is in a position to know that *q*. Since we know that *having hands* entails *not being a deceived bodiless brain in a vat*, if we cannot have reflective knowledge that we are not deceived bodiless brains in vats, then given PEC₁, we cannot have reflective knowledge that we have hands. The point can, of course, be generalized. Since radical skeptical hypotheses are incompatible with virtually all of the ordinary propositions we routinely take ourselves to know, if we lack reflective knowledge that radical skeptical hypotheses are false, then we lack reflective knowledge of the most mundane of ordinary propositions.

Pritchard contends that skeptical challenges force us to recognize the reflectively lucky nature of our anti-skeptical beliefs and that this, in turn, explains the enduring epistemic angst that skeptical hypotheses engender. Pritchard argues that the ineliminability of reflective luck shows that we not only lack reflective knowledge that the skeptic's hypotheses are false, we also lack reflective knowledge that our ordinary commonsense beliefs are true. If Pritchard is right, we may, indeed, possess a great deal of ordinary knowledge, but the ineliminability of reflective luck will forever preclude us from reflectively being able to tell that we do.

4. Conclusion

Reflecting on the nature and scope of epistemic luck gives us deeper insight into the nature and scope of knowledge. Gettier cases demonstrate that fallible justification is not capable of ruling out all forms of knowledge-destroying epistemic luck and that thus knowledge requires more than justified true belief. Just what anti-luck condition must be added justified true belief to arrive at an adequate analysis of knowledge remains an open question.

Recognizing which forms of epistemic luck are incompatible with knowledge and which are not puts us one step closer to identifying the correct luck-eliminating condition. It is now generally acknowledged that veritic luck is incompatible with knowledge. Whether other forms of epistemic luck, such as, justification-oriented luck, are incompatible with knowledge is a question that deserves more attention. At a minimum, any adequate theory of knowledge must be capable of ruling out all cases of veritic luck and to date no theory has been able to do so.

The possibility of knowledge-destroying veritic luck poses no special skeptical threat where first-order knowledge is concerned. As long as a person is not veritically lucky with respect to p , she will know that p , provided she has met the other conditions required for knowledge. The situation appears to be different where second-order knowledge is concerned. While there is no consensus to date as to how serious an obstacle the Gettier problem poses for second-order knowledge, it poses enough of an obstacle to such knowledge to render implausible the once widely held KK-thesis according to which knowing entails knowing that one knows.

Veritic luck is not the only form of epistemic luck that threatens more reflective forms of knowledge. Reflective luck also threatens to undermine the possibility of reflectively knowing that one knows. Our apparent inability to know, on the basis of reflection alone, that the skeptic's radical hypotheses are false, together with the principle of epistemic closure, threatens to undermine the possibility of reflective knowledge altogether.

5. References and Further Reading

Chisholm, Roderick. 1986. "The Place of Epistemic Justification." *Philosophical Topics* 14: 85-92.

Argues that one cannot generally know that one knows on the grounds that one cannot generally know whether or not one's evidence for p is defeated by Gettier considerations.

Chisholm, Roderick. 1964. "The Ethics of Requirement." *American Philosophical Quarterly* 1: 147-153.

Provides a *No Defeaters* response to the Gettier problem.

Clarke, Michael. 1963. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper." *Analysis* XXIV: 46-48.

Argues that *No False Grounds* is mistaken since S can derive a justified true belief that p from a justified true belief that q and still fail to know that p because S 's grounds for q are false. Contends that knowledge is "fully grounded" justified true belief, where in order to be *fully grounded*, the chain of reasons leading up to S 's proximate grounds for p must itself contain no false

grounds at any point in the chain.

Coffman, E.J. 2007. "Thinking about Luck." *Synthese* 158: 385-398.

Defends a *lack of control* account of luck.

Dretske, Fred. 1971. "Conclusive Reasons." *Australasian Journal of Philosophy* 49: 1-22.

Argues that in order to rule out knowledge-destroying luck, one's reasons for p must be conclusive in the sense that one would not have had those reasons if p were false.

Engel Jr., Mylan. 2000. "Internalism, the Gettier Problem, and Metaepistemological Skepticism." *Grazer Philosophische Studien* 60: 99-117.

Argues that the Gettier problem poses three distinct challenges to second-order knowledge which, when taken together, threaten to undermine the possibility of knowing that one knows.

Engel Jr., Mylan. 1992. "Is Epistemic Luck Compatible with Knowledge?" *Southern Journal of Philosophy* 30: 59-75.

Identifies veritic luck as the principal form of knowledge-destroying luck. Distinguishes veritic luck from evidential luck. Argues that, of these two types of luck, only veritic luck is incompatible with knowledge. Further argues that only externalist epistemologies are capable of ruling out veritic luck.

Feldman, Richard. 1981. "Fallibilism and Knowing that One Knows." *Philosophical Review* XC: 266-282.

Defends the iterative KK-thesis. Argues that the Gettier problem poses a minor, but hardly insurmountable, obstacle to second-order knowledge.

Feldman, Richard. 1974. "An Alleged Defect in Gettier Counterexamples." *Australasian Journal of Philosophy* 52: 68-69.

Provides a decisive example of an *all-true-evidence* Gettier case that shows that *No False Grounds* is too weak.

Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MA: Rowman and Littlefield Publishers.

Argues that second-order knowledge is too easy on externalistic accounts of knowledge and that, therefore, such accounts fail to capture the kind of knowledge that interests us.

Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121-3.

Demonstrates that justified true belief is not sufficient for knowledge. Highlights two paradigm examples of knowledge-destroying epistemic luck.

Greco, John. 2004. "A Different Sort of Contextualism." *Erkenntnis* 61: 383-400.

Defends a contextualist virtue epistemology. Offers a situationalist account of veritic luck, that is, an account tied to one's *epistemic situation* rather than to one's evidence: S is veritically lucky in believing that p if and only if, *given S 's epistemic situation*, it is just a matter of luck that S 's belief that p is true.

Greco, John. 2003. "Virtue and Luck, Epistemic and Otherwise." *Metaphilosophy* 34: 353-366.

Defends a virtue theoretic solution to the Gettier problem. Argues that when S has a true belief that p *because* S believes out of intellectual virtue (that is, when S 's believing out of intellectual virtue is what accounts for her having a true belief that p rather than a false belief or no belief), then S 's true belief that p is not veritically lucky.

Goldman, Alvin. 1979. "What Is Justified Belief?" In *Justification and Knowledge*. Ed. George Pappas. Dordrecht, Holland: D. Reidel Publishing Company.

Develops and defends an externalistic, process reliabilist account of justified belief.

Goldman, Alvin. 1967. "A Causal Theory of Knowing." *The Journal of Philosophy* 64: 355-372.

Attempts to solve the Gettier problem by replacing the traditional justification condition with a causal constraint requiring that one's belief that p be caused by the fact that makes p true.

Harman, Gilbert. 1973. *Thought*. Princeton, NJ: Princeton University Press.

Presents three much-discussed examples intended to show that knowledge can be undermined by readily available misleading evidence that one does not possess. Defends a *No Essential False Grounds* response to the Gettier problem.

Harper, William. 1996. "Knowledge and Luck." *Southern Journal of Philosophy* 34: 273-283.

Demonstrates that the Gettier problem plagues all fallibilistic theories of justification, both internalistic and externalistic alike. Argues that in addition to the traditional justification, truth, and belief conditions, an adequate analysis of knowledge must incorporate a "no luck" condition.

Heller, Mark. 1999. "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives*, 13, *Epistemology*: 115-29.

Proposes a context-sensitive modal account of epistemic luck according to which 'S's belief that p is epistemically lucky' is true if and only if there is at least one world (in a contextually-determined set of epistemically relevant worlds) where S's belief that p is false. Defends a contextualist anti-luck epistemology which maintains that 'S knows that p ' is true if and only if there is no world (in a contextually-determined set of epistemically relevant worlds) where S's belief that p is false.

Hetherington, Stephen. 2005. "Gettier Problems." *Internet Encyclopedia of Philosophy*.

Canvasses various purported solutions to the Gettier problem. Concludes with the contentious suggestion that justified true belief *is* sufficient for knowledge and that veritically-lucky justified true beliefs, like those in Gettier's original examples, are actually cases of knowledge.

Hiller, Avram and Ram Neta. 2007. "Safety and Epistemic Luck." *Synthese* 158: 303-13.

Demonstrates that *safe true belief* is not sufficient for knowledge by providing an example of a veritically lucky safe true belief that clearly falls short of knowledge.

Klein, Peter. 2008. "Useful False Beliefs." In *Epistemology: New Essays*. Ed. Quentin Smith. Oxford: Oxford University Press.

Argues that there can be *beneficial falsehoods*—falsehoods essential to one's justification—that nevertheless give one knowledge.

Klein, Peter. 1971. "A Proposed Definition of Propositional Knowledge." *Journal of Philosophy* 68: 471-482.

Defends a *No Defeaters* solution to the Gettier problem.

Lackey, Jennifer. 2007. "Why We Don't Deserve Credit for Everything We Know." *Synthese* 158: 345-361.

Argues that knowledge is not credit-worthy true belief and that thus virtue theoretic accounts of knowledge are mistaken.

Lackey, Jennifer. 2006. "Pritchard's Epistemic Luck." *The Philosophical Quarterly* 56: 284-289.

Provides a counterexample to Pritchard's modal account of luck.

Lehrer, Keith. 1990. *Theory of Knowledge*. Boulder, CO: Westview Press.

Defends a coherence theory of justification. Argues that knowledge is undefeated justified true acceptance.

Lehrer, Keith and Thomas Paxson Jr. 1969. "Knowledge: Undefeated Justified True Belief." *The Journal of Philosophy* 66: 225-237.

Defends an alternative *No Defeaters* response to the Gettier problem. Introduces the Tom Grabit counterexample to the Chisholm/Klein account of defeaters.

Lycan, William. 1977. "Evidence One Does Not Possess." *Australasian Journal of Philosophy* 55: 114-26.

Argues that misleading evidence one does not possess does not undermine one's knowledge and that, thus, Harman Cases are actually cases of knowledge.

Murphy, Peter. 2005. "Closure Failure for Safety." *Philosophia* 33: 331-34.

Adapts Kripke's famous "blue barn" counterexample to Nozick's analysis of knowledge to show that *safe true belief* accounts of knowledge also result in closure failure.

Myers, Robert G. and Kenneth Stern. 1973. "Knowledge without Paradox." *Journal of Philosophy* 70: 147-160.

Argues that p can justify S in believing some other proposition q only if p is true and that, thus, Gettier's original cases do not provide examples of *justified* true beliefs that fall short of knowledge.

Nozick, Robert. 1981. "Knowledge and Skepticism." In *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Develops and defends a sensitivity-based subjunctive conditionals analysis of knowledge.

Plato: *Theaetetus* in *Plato: Collected Dialogues*. Eds. E. Hamilton and H. Cairns. Princeton, NJ: Princeton University Press, 1961.

Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Oxford University Press.

Proposes a modal account of veritic luck. Argues that *safety* precludes veritic luck. Defends an externalist neo-Moorean *safe true belief* account of ordinary knowledge. Concedes to the skeptic that *reflective luck* is ineliminable and that such luck is incompatible with reflective knowledge.

Pritchard, Duncan. 2003. "Virtue Epistemology and Epistemic Luck." *Metaphilosophy* 34: 106-30.

Pritchard, H.A. 1950. *Knowledge and Perception*. Oxford: The Clarendon Press.

Riggs, Wayne. 2007. "Why Epistemologists Are So Down on their Luck." *Synthese* 158: 329-344.

Defends a *lack of control* account of luck. Argues that knowledge is credit-worthy true belief.

Roth, Michael. 1990. "The Wall and the Shield: K-K Reconsidered." *Philosophical Studies* 59: 147-157.

Argues that the Gettier problem poses no obstacle to second-order knowledge on the grounds that Gettier-type contingencies lie safely outside the wall of fallibilism and can simply be ignored (unless they actually obtain).

Russell, Bertrand. 1912. *The Problems of Philosophy*. Oxford, England: Oxford University Press.

Skyrms, Brian. 1967. "The Explication of 'X Knows that p'." *Journal of Philosophy* 64: 373-389.

Provides one of the first cases of an *all true evidence* Gettier case.

Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Volume 1. Oxford, England: Oxford University Press.

Develops and defends a virtue epistemology which maintains that *knowledge is reliably-produced safe true belief, the correctness of which is attributable to one's epistemic competence*. Argues that when the correctness of a reliably-produced safe belief is attributable to the proper exercise of an epistemic competence, then the resultant belief is not epistemically lucky.

Sosa, Ernest. 2000. "Skepticism and Contextualism." *Philosophical Issues*, 10, *Skepticism*: 1-18.

Defends a non-contextualist safety-based Moorean response to the skeptical paradox.

Sosa, Ernest. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives*, 13, *Epistemology*: 141-54.

Develops a safety-based Moorean response to the skeptical paradox. Argues that such a response is preferable to skeptical, tracking, relevant-alternative, and contextualist accounts.

Steup, Matthias. 2006. "The Analysis of Knowledge." *Stanford Encyclopedia of Philosophy*.

Discusses the necessary conditions for knowledge. Examines internalistic and externalistic analyses of knowledge.

Swain, Marshall. 1981. *Reasons and Knowledge*. Ithaca, NY: Cornell University Press.

Truncellito, David A. 2007. "Epistemology," *Internet Encyclopedia of Philosophy*.

Unger, Peter. 1968. "An Analysis of Factual Knowledge." *Journal of Philosophy* 65: 157-70.

Shows that various forms of epistemic luck – including propositional luck, existential luck, and facultative luck – are compatible with knowledge. Argues that knowledge is non-accidentally true belief.

Vahid, Hamid. 2001. "Knowledge and Varieties of Epistemic Luck." *Dialectica* 55: 351-362.

Argues that truth-oriented veritic luck and justification-oriented luck are both incompatible with knowledge.

Vogel, Jonathan. 1999. "The New Relevant Alternatives Theory." *Philosophical Perspectives*, 13, *Epistemology*: 155-80.

Uses *Hole-in-One* to argue that sensitivity is not necessary for knowledge.

Williamson, Timothy. 2000a. "Scepticism and Evidence." *Philosophy and Phenomenal Research* 60: 613-28.

Williamson, Timothy. 2000b. *Knowledge and Its Limits*. Oxford, England: Oxford University Press.

Zagzebski, Linda. 1996. *Virtues of the Mind*. Cambridge, England: Cambridge University Press.

Characterizes Gettier-style knowledge-destroying luck as cases of “double luck” where epistemic bad luck is cancelled out by epistemic good luck. Argues that no fallibilist epistemology can rule out knowledge-destroying luck. Defends a virtue-based epistemology according to which *knowledge is a state of cognitive contact with reality arising out of acts of intellectual virtue*, and argues that this definition of knowledge is immune to the Gettier problem because truth is entailed by the other components of the definition.

Author Information

Mylan Engel Jr.

Email: mylan-engel@niu.edu

Northern Illinois University

U. S. A.