

PART IV

Gettier and Inferential Knowledge

16

Inferential Knowledge and the Gettier Conjecture

Rodrigo Borges

Abstract

The Gettier Problem, when properly understood, has a straightforward solution. My main thesis—my ‘Gettier conjecture’—is that gettierized subjects fail to know in virtue of their justified true belief depending causally and evidentially on something they fail to know. Inferential knowledge (i.e., knowledge of a conclusion) requires knowledge of all the premises on which one’s conclusion depends causally and evidentially. This chapter is a first effort in trying to support this thesis. Section 2 discusses the Gettier conjecture, the notions of evidential and causal dependence, and applies the conjecture to the original Gettier cases. Section 3 looks at two objections: the claims that the conjecture fails to deal with all Gettier cases and that the Gettier Problem is a philosophical ‘dead end.’ In Section 4 I offer further support for the conjecture by situating it within a knowledge-first framework.

1 Introduction

The Gettier Problem, when properly understood, has a straightforward solution. My main thesis—my ‘Gettier conjecture’—is that gettierized subjects fail to know in virtue of their justified true belief (JTB) depending causally and evidentially on something

they fail to know. Inferential knowledge (i.e., knowledge of a conclusion) requires knowledge of all the premises one's conclusion depends causally and evidentially on. This chapter makes a case for this thesis. To that end, the discussion below is structured in the following way. Section 2 discusses the Gettier conjecture, the notions of evidential and causal dependence, and applies the conjecture to the original Gettier cases. Section 3 looks at two objections: the claims that the conjecture fails to deal with all Gettier cases and that the Gettier Problem is a philosophical 'dead end.' In Section 4 I offer further support for the conjecture by situating it within a knowledge-first framework.

2 The Gettier Conjecture

Before we go any further, let me state our conjecture more explicitly:

(GC) Necessarily, the protagonist of a Gettier case fails to know that *p* (even though she has a justified true belief that *p*) in virtue of *p* depending on at least one premise she does not know.

To see how the Gettier conjecture works, consider Gettier's Case I.¹ Smith has a justified true belief that the man who will get the job has ten coins in his pocket. He infers this proposition from his justified but false belief that Jones will get the job and Jones has ten coins in his pocket. Smith does not know that the man who will get the job has ten coins in his pocket. But why? According to the Gettier conjecture, Smith fails to know the conclusion of his inference because his belief in the conclusion depends on a false (and, thus, unknown) premise.

GC also applies straightforwardly to Gettier's Case II.² Smith arrives at the true and justified belief that either Jones owns a Ford or Brown is in Barcelona, via inference, from his justified but false belief that Jones owns a Ford. Smith does not know the conclusion of his inference. Again, if GC is true, then the reason why Smith fails to know is because he fails to know the premise his conclusion depends on.

GC commits us to the claim that *all* Gettier cases involve inference. It commits to the claim that all gettierized beliefs are inferential beliefs. Why think that? At this point, let me give a few reasons to believe that this is the right thing to say about Gettier cases in general. I will have something to say about why denying this claim leads to problems in Section 2.1. For now, let me focus on a positive reason to think why this is right.

Gettier not only presented two cases which show that knowledge is not the same as justified true belief, he also gave us a blueprint of those cases. He gave us a set of instructions we can use to multiply them. The instructions were something like 'In order to produce a justified true belief that is not a case of knowledge, have the protagonist of your case competently deduce a true conclusion from a justified but false premise.'

¹ Gettier (1963: 122).

² Gettier (1963: 122–3).

Cases I and II not only refuted the JTB account of knowledge, the cases also highlighted an absurd consequence of this account. If the JTB account were correct, then competent deduction from justified false beliefs would always yield knowledge of its conclusion. This does not hold true of Cases I and II, however. To neglect the generality of Gettier's case against the JTB account of knowledge is to neglect an important epistemological lesson.

Gettier's blueprint was laid out in the form of two principles about justification he asked his readers to accept.

First, in that sense of 'justified' in which S's being justified in believing P is a necessary condition of S's knowing that P, it is possible for a person to be justified in believing a proposition that is in fact false. Second, for any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q.³

Call the first principle about justification 'Fallibility' (there are justified *false* beliefs) and the second 'Justification Closure' (justification is closed under known entailment). According to our interpretation of Gettier's blueprint for his cases, a case C is a Gettier case only if Fallibility and Justification Closure are instantiated in C. So, the fact that those principles are instantiated is a *necessary* condition on a case being a Gettier case—i.e., a necessary condition a case has to satisfy in order for it to qualify as a case of the type Gettier was talking about.⁴

Of course, we might also ask if satisfying these principles is also sufficient to make a case a Gettier case. It seems like it is not and looking at an example will help bring this out. The case below will also help bring out another important aspect of GC—the notion of dependence.⁵ Consider the following case:⁶

The Harmless Ignorance Case

Smith arrives through reasoning to the belief that (p) someone in his office owns a Ford. His premises are (q) that Havit owns a Ford, (r) that Havit works in his office, (s) that Nogot owns a Ford, and (t) that Nogot works in his office. As it turns out s is false while q, r, and t are true.

Most people think Smith knows that someone in his office owns a Ford. I agree that Smith knows in this case. But the case instantiates both Fallibility (Smith is justified in believing *falsely* that Nogot owns a Ford) and Justification Closure (Smith *competently deduces* that someone in Smith's office owns a Ford from his justified false belief that Nogot owns a Ford and from his true belief that Nogot works in his office). The upshot

³ Gettier (1963: 121). ⁴ Shope (1983: 4) concurs.

⁵ If the so-called 'knowledge from non-knowledge' cases discussed by Klein (2008), Warfield (2005), and others in fact showed that knowledge may depend on non-knowledge (in the sense of 'depend' to be defined momentarily), then those cases would also show that instantiating Fallibility and Justification Closure is not sufficient to turn a case into a Gettier case, for both principles are instantiated in alleged cases of knowledge from non-knowledge. I discuss those cases and why they do not show that knowledge may depend on non-knowledge in Section 3.

⁶ Adapted from Lehrer (1965).

is that instantiating Fallibility and Justification Closure is *not* sufficient to turn a case C into a Gettier case.

We can use the notion of dependence that appears in GC to capture the difference between cases in which the instantiation of Fallibility and Justification Closure prevent the subject from knowing and cases—like the Harmless Ignorance case—in which the instantiation of those principles does not prevent the subject from knowing. Here is what I take the notion of dependence in GC to imply:

(Dependence) *p* depends on *q* for *S* in a case *C* iff (i) *S* would not have believed that *p* had she not believed that *q* in *C*; and (ii) if *q* were not in *S*'s evidence set in *C*, *p* would not have been justified for/known by *S* in *C*.⁷

There might be a sense in which *p* depends on *q* if either (i) or (ii)—but not both—are satisfied, but that is not the sense of 'p depends on q' that I am interested in here. Furthermore, while (i) captures the sense in which *p* depends *causally* on *q* in a crucial way,⁸ (ii) captures the sense in which *p* depends *evidentially* on *q* in a crucial way. I will thus refer to (i) and (ii) as the *causal counterfactual* and the *evidential counterfactual*, respectively.

Now, going back to the distinction between Gettier cases and the Harmless Ignorance case, we can now state precisely what the difference between those cases is: even though Fallibility and Justification Closure are instantiated in both cases, Dependence is satisfied by Gettier cases but not by the Harmless Ignorance case.

A word of caution. One could plausibly object that in Harmless Ignorance, Smith's true belief that someone in his office owns a Ford is *causally overdetermined* by his false and true beliefs about who owns a Ford. When we apply the causal counterfactual in Dependence to this case, the result seems to be that neither the true nor the false belief is the cause of Smith's true belief in the conclusion, for, had Smith not believed either one of them, the other would have been sufficient to bring about Smith's belief in the conclusion. Worse still, once we make this problem salient, one starts to worry whether the causal counterfactual will infect Dependence with the disease of deviant causation that plagues analyses of *causation*. In fact, one may argue that it is only a matter of time until we are presented with cases featuring other types of deviant causal chains (e.g., cases of early and late preemption). The objector would then point out that those cases all falsify Dependence.

These worries are misplaced and emerge from a misunderstanding of the role Dependence plays in my story. Dependence is a heuristic we apply to determine if a proposition is causally and/or evidentially crucial to another proposition (for a

⁷ This account of dependence is similar to the one proposed by Klein (2008).

⁸ This kind of counterfactual dependence holds even if causation cannot be reduced to counterfactuals as in Lewis (1979); and even if counterfactuals cannot be reduced to causation as in Jackson (1977). The idea that beliefs are causes (in particular, of other beliefs) is widely accepted. But for some dissenting arguments see Klein (2012) and McLaughlin (2006). For a recent synoptic discussion of the relationship between causation and counterfactual conditionals, see the excellent Paul & Hall (2013) and the references therein.

particular subject, in a particular case). Dependence is *not* a necessary *or* sufficient condition on belief causation (or on causation in general, for that matter). Thus, Dependence is not really infected with the issues plaguing analyses of *causation*.

In sum, Dependence *does not* uncover what the *true cause* of a belief is, it tells us only whether one's conclusion depends—in an epistemically significant way—on another item in one's evidence set.⁹

2.1 Objection: Gettier cases without Gettier's principles

Above I defended the idea that, necessarily, if a case C is a Gettier case, then C instantiates Fallibility and Justification Closure. Many implicitly or explicitly deny this claim. Those philosophers take cases in which either Fallibility or Justification Closure (or both) are not instantiated to be Gettier cases. If they are right, then GC does not apply to those cases and is, thus, not a fully general explanation of the Gettier phenomenon.

If GC is true, then all Gettier cases are cases in which the subject's justified true belief is the result of a deductively valid inference. However, some alleged Gettier cases do not involve any inference. Alvin Goldman's barn-facade case¹⁰ is sometimes taken to be a Gettier case in which the protagonist acquires a justified true belief via perception. Bertrand Russell's stopped-clock case¹¹ and Roderick Chisholm's sheep-in-the-field case¹² are also sometimes taken to be Gettier cases whose protagonists acquire a justified true belief via perception.¹³ If those are Gettier cases in which the target justified true belief is not acquired through inference, then we could not use GC to explain why their protagonist fails to know.

There are a few things we can say in response to this worry.

First, it is not clear that the protagonist in all of those cases fails to know. Second, even if their protagonist fails to know, a case may be made that her justified true belief depends on unknown propositions. Third, the view that takes those cases to be Gettier cases runs the risk of overgenerating Gettier cases. I will discuss each of these issues in turn.

Ernest Sosa has famously argued that there is an important sense in which the protagonist of the barn-facade case knows that that is a barn—the *animal knowledge* sense of knowledge—while there is another sense in which he fails to know that that is a barn—the *reflective knowledge* sense of knowledge.¹⁴ William Lycan also reports not

⁹ I will have more to say about Dependence in Section 3, when I discuss the more general principle GC is an instance of—the knowledge-from-knowledge principle.

¹⁰ Cf. Goldman (1976: 772–3).

¹¹ Russell (1948/2009).

¹² Cf. Chisholm (1977).

¹³ Russell originally used the stopped-clock example to show that there can be true belief without knowledge (1948/2009: 140). Israel Scheffler (1965: 112) was the first to suggest that the case could be turned into a counterexample to the JTB analysis of *knowledge*, if we assume that the subject looking at the clock 'has good grounds to suppose the clock *is* going.' This formulation of Russell's case turns it into a case in which the protagonist's justified true belief that it is 10:00 a.m. depends on the falsehood 'the clock *is* going.' I come back to the issue of whether or not this is a non-inferential case below.

¹⁴ Cf. Sosa (2007: 96 fn. 1) and Sosa (2011: 82–95). In Sosa (2011: 92), Sosa tries to accommodate the fact that "many of us cannot believe that the fake barns subject knows at any level whatsoever, whether animal or reflective" by distinguishing still a further sense in which one can be truly said to know something—i.e.,

having the intuition that the person in this case fails to know.¹⁵ Finally, Tamar Gendler and John Hawthorne have challenged the very reliability of intuitive judgments in response to barn-facade cases.¹⁶

The point of mentioning what these philosophers have to say about barn-facade cases is that there is wide disagreement as to whether the protagonist in Goldman's case knows. This contrasts sharply with the fact that the wide majority of epistemologists agree that the protagonist in the original Gettier cases fails to know.¹⁷

However, even if the protagonist of those non-inferential cases fails to know, it can still be argued that their justified true belief depends on propositions they fail to know. For example, one could insist that Russell's justified true belief that it is 10:00 a.m. depends on the false proposition that the clock is working. Similarly, one may insist that Chisholm's justified true belief that there is a sheep in the field depends on the false proposition that that is a sheep, or that Goldman's justified true belief that that's a barn depends on the false proposition that there are no barn look-a-likes around.¹⁸ Under this interpretation these *are* Gettier cases, since they too instantiate Fallibility and Justification Closure.

Ultimately, we cannot settle for this approach to those cases. There is still an important difference between those cases and the original Gettier cases: while the protagonist of the original Gettier cases *explicitly deduces* his true belief from the false one, the protagonist of the cases we are considering does no such thing.

The defender of GC can point out that the view we are considering faces another, more fundamental problem. The view denies that Gettier cases necessarily instantiate Fallibility or Justification Closure (or both). Instead, the view argues that there being a justified true belief that is not a case of knowledge is what makes a case a Gettier case. The problem is that this suggestion is certainly false. Many cases in which the protagonist has a justified true belief that is not a case of knowledge are also *not* Gettier cases. Even if true and justified, my belief that my ticket is lost is not a case of knowledge.¹⁹ Most importantly, this is also not a case of gettiered belief. The view that says that we can have Gettier cases without Gettier's principles is too permissive; it takes cases that are clearly not Gettier cases to be Gettier cases. The view I proposed does not have this problem. The cases discussed by Russell, Goldman, and Chisholm feature a justified true belief that is not knowledge, but they are not Gettier cases.²⁰

the *human knowledge* sense of knowledge. Human knowledge, according to Sosa, comes in degrees and in its lowest degrees it corresponds to reflective knowledge; in its highest degree, human knowledge involves "scientific and even philosophical perspectives that enable defense of one's first-order belief as apt." Sosa argues that those who think that Henry in the barn-facade case doesn't know, full stop, have the *human knowledge* sense of knowledge in mind.

¹⁵ Cf. Lycan (2006: 158, 162–3).

¹⁶ Gendler & Hawthorne (2005).

¹⁷ Weatherson (2003) and Hetherington (2011) are notable exceptions.

¹⁸ The reader can no doubt work out the truth of the relevant evidential and causal counterfactuals.

¹⁹ Here, I am supposing that I form the belief before I learn about the lottery results.

²⁰ I include Richard Feldman's Clever Reasoner case in this larger class of cases that are not Gettier cases but in which the protagonist has a justified true belief that is not a case of knowledge. See Feldman (1974).

2.2 Objection: the Gettier Problem is a philosophical dead end

In Section 3 I will be launching an argument in support of GC from a knowledge-first epistemology perspective. However, Timothy Williamson, the main proponent of the knowledge-first epistemology program, is usually taken (with justification) to be hostile to the idea that we epistemologists should be spending time trying to solve the Gettier Problem. If I take my view to be (broadly speaking) a knowledge-first view, then why do I care about the Gettier Problem?

First of all, what do I mean by ‘knowledge-first epistemology’? I will not engage in Williamsonian exegesis here. For me, knowledge-first epistemology takes knowledge to be explanatorily prior to belief.²¹ For example, I, like Williamson, think that epistemological normativity should be explained in terms of knowledge norms rather than belief or justified belief norms.

When it comes to the Gettier Problem, ultimately I want to explain why justified true belief in Gettier cases fails to be knowledge by appealing to the fact that one has inferential knowledge only if one knows all the propositions this knowledge depends on. This will become clear in Section 3.

We can say something a little less programmatic about how knowledge-first epistemologists should think about the Gettier Problem, however. For instance, we might distinguish between two (perhaps related) *Gettier problems*.

As usually understood, ‘Gettier Problem’ refers to a problem in conceptual analysis. Gettier (and virtually everyone since his paper was published) thought his cases refuted the conceptual claim that *knowledge* and *justified true belief* are necessarily co-extensive. In the two cases discussed by Gettier, the protagonist instantiates the concept *justified true belief* but not the concept *knowledge*. Call this the *Conceptual Gettier Problem*.

Some epistemologists trying to solve the *Conceptual Gettier Problem* have argued that a fourth condition must be added to the traditional analysis.²² Others have proposed analyses of *knowledge* that reject the traditionalist claim that justification is a necessary condition on *knowledge* and put forward a counterfactual condition instead.²³

In *Knowledge and Its Limits* Williamson pointed out that the (then) forty-something years of work on the Conceptual Gettier Problem had yielded no consensus about how to define *knowledge*. Williamson takes this history of failed analysis to provide inductive evidence that a non-circular analysis of *knowledge* is not forthcoming.²⁴

As many have pointed out, Williamson’s argument is open to the objection that most, if not all, philosophical concepts (e.g., *cause*, *justice*, etc.) have resisted analysis for millennia, but this never stopped philosophers from trying to define these concepts;

²¹ Cf. Williamson (2011).

²² For example, Klein (1971), Chisholm (1977), Goldman (1986), Lehrer (1990), Feldman (2002), and Pritchard (2012), to name only a few.

²³ For example, Nozick (1981), and Dretske (1971).

²⁴ Williamson (2000: 30).

if for no other reason because doing so significantly deepens our understanding of a wide range of other philosophically interesting concepts (e.g., *similarity*, *fairness*, etc.).

I will not try to resolve the dispute between Williamson and his critics.²⁵ Rather, I want to point out that both parties take the Gettier Problem to be a problem in conceptual analysis and that, as a result, whether this problem is interesting, relevant, or worth thinking about, depends in part on whether or not we can offer a non-circular definition of *knowledge*. Be that as it may, I do not think that Gettier cases pose only (or even primarily) a conceptual problem. That is, I do not think that that class of cases is epistemologically interesting only because it challenges us to come up with a non-circular definition of *knowledge*. Gettier cases pose an epistemologically important problem whether or not it is possible for us to give a non-circular analysis of *knowledge*. In particular, Gettier cases pose an *explanatory challenge*: what exactly explains why protagonists in Gettier cases fail to know? We know that the protagonist of those cases have a true belief that is justified (in some sense of ‘justified’). But what prevents their true belief to be the kind of true belief we call ‘knowledge’? Call this challenge raised by Gettier cases the *Explanatory Gettier Problem*.

We do not need to solve the Conceptual Gettier Problem or decide if *knowledge* can receive a non-circular analysis before we tackle the Explanatory Gettier Problem. We can approach this explanatory problem raised by Gettier without dwelling on the nature of *knowledge*.

But is tackling the Explanatory Gettier Problem worth our time? I think it is, and, like the proverbial pudding, the proof of this claim is in the ‘eating.’ Its value emerges from the theoretical fruits it bears and from its overall fitness with other epistemological views. The Gettier Conjecture is an answer to the explanatory challenge posed by Gettier cases and we have already seen that this principle bears some important epistemological fruits: it forces us to get clear on the notion of causal and evidential dependence and it forces us to offer a more principled account of what makes a case a Gettier case. In Section 3 I will show how this conjecture fits a broader epistemological outlook.

3 The Knowledge-from-Knowledge Principle

Is there any reason for us to accept the Gettier Conjecture besides what it says about Gettier cases? In this final section I show that there is a lot more we can say in favor of GC. GC follows from a more general principle about inferential knowledge. This general principle was explicitly and implicitly accepted by most philosophers before Gettier. The principle is also a consequence of plausible accounts of evidence and evidential defeat. And, finally, depending on how we understand justification, a version of the

²⁵ Williamson (2000: 33–4) has at least one other argument against the claim that *knowledge* can be given a non-circular definition. There are reasons to think that the most popular conditions on knowledge proposed by epistemologists should be themselves understood in terms of *knowledge*. Perhaps, it is argued, *justification* (Sutton 2007), *safety* (Williamson 2000: chapter 7), or even *belief* (Williamson 2011) are definable only by appeal to *knowledge*.

no-false-grounds solution to the Conceptual Gettier Problem may be seen as a notational variant of this more general principle behind GC.

GC is a special case of the more general principle stating that inferential knowledge (i.e., knowledge of the conclusion of an inference) requires knowledge of all the propositions that conclusion depends on. Call this general principle the *Knowledge-from-Knowledge Principle*:

(KFK) S has inferential knowledge that p only if S knows all propositions on which p depends.

KFK provides a basis for the intuitive notions that reasoning generates knowledge only if it is 'sound,' and that one knows the conclusion of one's reasoning only if one's premises provide one with a 'good reason' to accept it. It also explains why criticisms of the form 'You should not believe that—the truth of that depends on something you don't know' seem appropriate, when true. If this thesis is true and one acquires inferential knowledge only if one knows all propositions on which one's conclusion depends on, then our preference for sound reasoning over unsound reasoning is in part explained by our preference for known premises over premises we fail to know, together with the fact that we assign great epistemic value to knowledge. Furthermore, this claim about inferential knowledge has been widely accepted throughout the history of philosophy.

For Aristotle, one 'demonstrates' that p is true only if one has a syllogistic argument for p and one's premises are items of scientific knowledge ('primitives').²⁶ René Descartes and John Locke espoused similar views. Consider this passage from Descartes' *Rules for the Direction of the Mind*: "many facts which are not self-evident are known with certainty, provided they are inferred from true and known principles through a continuous and uninterrupted movement of thought in which each individual proposition is clearly intuited."²⁷ Like Aristotle, Immanuel Kant also took inferential knowledge to be the business of syllogistic demonstrations. For Kant the major premise in a syllogism is an instance of 'universal *a priori* knowledge,' a principle of considerable generality. It follows that, for Kant, every syllogism is a mode of deducing knowledge from a principle.²⁸ But, since, for Kant, something is a principle just in case it is an item of *a priori* knowledge, the claim that every syllogism is a mode of deducing knowledge from a principle entails that one's inference produces knowledge only if one is certain (and, therefore, knows) that all the premises in one's argument are true.²⁹

²⁶ Aristotle says that, given this account of demonstration, "if we know and are convinced of something because of the primitives, then we know and are convinced of them better, since it is because of them that we know and are convinced of the posterior items" (Aristotle 1994: 72a25–30). Jonathan Barnes' commentary on this passage in Aristotle (1994: 101–2) confirms our reading of the Aristotelian passage. Note also that, while Aristotle is claiming that we 'know better' the conclusions based on primitives than we know the primitives themselves, KFK makes no such commitment. KFK is silent on the issue of whether inference enhances the epistemic status of conclusions or not.

²⁷ Descartes (1985: 15). According to Locke, "every step in Reasoning, that produces Knowledge, has intuitive certainty" (Locke 1975: BK IV, chII, sec.7).

²⁸ Kant (1950: 301).

²⁹ Of course, it is not enough that one knows the major premise in one's syllogism; one also has to know the minor premise of one's syllogism, but this is assumed by all the people we are discussing.

In the twentieth century Bertrand Russell echoed the tradition's tacit acceptance of KFK in the classic *The Problems of Philosophy* (1912). While Aristotle, Descartes, Locke, and Kant seem to be particularly concerned with what we could call *scientific knowledge* (roughly, knowledge that entails certainty), Russell is concerned with knowledge in general: "But are we to say that nothing is knowledge except what is validly deduced from true premises? Obviously we cannot say this... In the first place... because it is not enough that our premises should be true, they must also be known."³⁰ This is a partial (but impressive) list of philosophers who have endorsed the thesis that inferential knowledge requires knowledge of all the premises one's conclusion depends on. Strictly speaking, with the exception of Russell, they all endorse something stronger than our thesis; something like the claim that *one knows the conclusion of one's reasoning only if one is certain of one's premises and one's premises entail one's conclusion*. We need not join these philosophers in thinking that only deduction from premises we are certain of yield inferential knowledge. KFK is compatible with strong inductive arguments from premises one knows, but of which one is not certain, also yielding inferential knowledge. For the philosophers I am discussing, on the other hand, inductive arguments only occasionally generate knowledge, most of the time they generate mere *probable opinion*. The fact remains that the stronger claim those philosophers are making directly entails KFK.³¹

If philosophical support for KFK was ubiquitous before Gettier, why hasn't anyone appealed to it in response to Gettier's cases? As a matter of fact, at least one influential epistemologist has reacted to Gettier's paper with an endorsement of KFK and GC. In his 1973 book *Belief, Truth and Knowledge* David Armstrong says the following about Gettier's original cases:

Gettier produces counterexamples to the thesis that justified true belief is knowledge by producing true beliefs based on justifiably believed grounds... but where these grounds are in fact

³⁰ Russell (1912: 132–3).

³¹ Things are a little more complicated than I am suggesting here. Aristotle, Descartes, Locke, and Kant seem to be talking about a necessary condition on inferential *scientia*, while Russell and I are talking about KFK, a necessary condition on inferential *cognitio*. In other words, after reading the passages I quoted above, one could claim that the view of inferential knowledge advanced by Aristotle, Descartes, Locke, and Kant does not entail the view Russell and I defend, because the view those philosophers are discussing is a view about *inferential scientia* and *scientia* is different in kind from *cognitio*—for one, the evidential standard for the former kind of state is a lot more demanding (i.e., certainty) than the evidential standard the latter kind of state has to meet. I do not think this is quite right. Even if *scientia* and *cognitio* are different (complex or simple) kinds of mental states, plausible bridge principles connecting the two can be easily devised. Those bridge principles would show that *scientia* does in fact entail *cognitio*. The entailment is just more indirect than I am suggesting in the body of the text. For example, if one has *scientia* that the conjunction of all the members in the set $\{q, q_1 \dots q_n\}$ is true, then one also has *cognitio* that it is true, for being certain entails both belief and an epistemic support sufficient for *cognitio*. Moreover, if $\{q, q_1 \dots q_n\}$ entails p and one acquires *scientia* that p via competent deduction, then one also acquires *cognitio* that p , for, again, *scientia* entails belief, the epistemic support required for *scientia* exceeds the one required for *cognitio*, and one's belief in p is at least as justified as one's belief in the least justified premise in $\{q, q_1 \dots q_n\}$. If this is right, then, even though a more complicated account of how Aristotle's view entails KFK is required, it is not something we cannot achieve. Thanks to Peter Klein for discussion here.

false. But because possession of such grounds could not constitute possession of knowledge, I should have thought it obvious that they are too weak to serve as suitable grounds.³²

This is a clear endorsement of KFK and of the Gettier Conjecture. I will not venture an explanation of why no one else in the post-Gettier literature reached for KFK in reaction to the Gettier cases. The most reasonable explanation of that fact probably involves a mixture of theoretical and sociological factors. Rather, my goal with this brisk historical contextualization of KFK is to show that the post-Gettier literature is the historical outlier when it comes to the way in which philosophers think about inferential knowledge. In general, the epistemological views of philosophers included at least a tacit commitment to KFK.

Of course, one might look at all this and think ‘So what? Philosophers in the past have been mistaken about all kinds of things. Everything you said is compatible with Gettier having uncovered the fact that philosophers have tacitly relied on something false, namely KFK.’ Fair enough. Nevertheless, we can do better than listing famous philosophers who accepted KFK. The rest of this section looks at direct arguments in support of KFK.

KFK receives support from our intuitive reaction to particular cases. Consider the following case modified from Vogel (1990). Suppose I ask you whether you will walk home after work and that this prompts you to reason thus:

1. My car is in the parking lot.
2. If my car is in the parking lot, then I will drive home.

Thus,

3. I will drive home.

Now, suppose that, on the basis of this argument, you tell me that you will not be walking home. We can now imagine two different versions of this case. In the first version, circumstances are as described and you know that your car is in the parking lot. It seems natural to say that, in this scenario, you know you will not be walking home.³³ Consider, however, a second version of this case. In this version I tell you that, lately, many cars have been stolen from the parking lot in which you parked your car. In this version of the case it seems that you do not know that your car is in the parking lot, since your evidence is not strong enough to rule out the hypothesis that your car has been stolen.

This second version of the case is the important one for us here. In this version, you not only fail to know that your car is in the parking lot but you also no longer know the conclusion of your inference (i.e., that you will not walk home). If you do not know

³² Armstrong (1973: 152).

³³ Of course, if my car is in the parking lot without gas or it has been smashed by a truck, I know my car is in the parking lot, but I do not know I will not have to walk home. I am, therefore, presupposing that we are talking about my *working car*.

that your car is in the parking lot and that is the only mode of transportation available to you, then you do not know that you will not have to walk home. But this is precisely the result we should expect if inferential knowledge required knowledge of all propositions your conclusion depends on. Notice also that in both versions we may assume that it is not only true that your car is in the parking lot, but also that this is very probable on your evidence. Still, this true and very probable belief is not knowledge in the second version of the case, and this prevents you from knowing that you will not be walking home.

A version of the lottery case can be used to support KFK as well. Suppose you bought a ticket in a large and fair lottery. The odds that any particular ticket will win are as abysmal as you wish (just make the lottery as large as you want). Suppose further that the drawing of the lottery took place at noon today and that it is now 1:00 p.m. Without yet having looked at the lottery result, you reason, at 1:15 p.m., in the following way:

1. My ticket lost.
2. If my ticket lost, then I will not be able to afford that Bob Dylan Stratocaster they will be auctioning at Christie's next week, for I have no other source of money I can use to buy it.

Therefore,

3. I will not be able to afford that Bob Dylan Stratocaster they will be auctioning at Christie's next week, for I have no other source of money I can use to buy it.

As with the second version of the previous case, you do not know the conclusion of your argument even if we suppose that it is true that you lost the lottery, and that you are basing your belief on your knowledge of the odds, making your belief highly likely (but not certain) to be true. Even though this claim is highly likely to be true, given your evidence, you do not know that your ticket lost. In this scenario—like in the second version of the car case above—you fail to know the conclusion of your inference and, again, this is exactly the result we should expect if KFK is true, for you fail to know one of the premises on which that conclusion depends.

What is more, if we change the lottery case so that you come to know the premise that your ticket lost *before* you infer that you will not be able to buy Dylan's Stratocaster, then it seems that you know the latter proposition.

So, KFK enjoys intuitive support from cases such as the car case and the lottery case. We can do even better than that, though. We can show that the knowledge-first account of evidence, Peter Klein's defeasibility theory of knowledge, and a plausible version of the no-false-grounds solution to the Conceptual Gettier Problem either entail or are compatible with KFK.

Consider the account of evidence according to which all and only knowledge is evidence (i.e., $E=K$).³⁴ Since, according to this view, evidence is what justifies

³⁴ Williamson (2000).

knowledge, it follows that knowledge is what justifies knowledge. So, only reasoning from what is known justifies knowledge. KFK is vindicated since $E=K$ entails that only known relevant premises justify knowledge. Even though $E=K$ entails KFK, the converse is not true. KFK merely requires that evidence relevant to one's conclusion be known, while $E=K$ requires that *all* evidence be known whether relevant or irrelevant to one's conclusion.

It can also be shown that a version of the *no-false-grounds* solution to the Conceptual Gettier Problem is a close cousin of KFK.

In an early reply to Gettier's paper, Michael Clark³⁵ proposed that the protagonist in Gettier cases fails to know because his justified true belief depends on at least one false belief. Alvin Goldman's account of inferential knowledge in *A Causal Theory of Knowing* also required that all beliefs involved in an inference be true if they are to yield knowledge of the reasoning's conclusion.³⁶ Gilbert Harman also defended a version of this view in several places.³⁷ More recently, Richard Feldman,³⁸ William Lycan,³⁹ and Richard Foley⁴⁰ have also endorsed similar views. The unifying idea behind all these *no-false-grounds* views (NFG) is that subjects in Gettier cases fail to acquire inferential knowledge because the conclusion of their inference depends on at least one false proposition.⁴¹ Those philosophers take NFG to be true of inference in general. Since knowledge entails true belief, whenever one fails to satisfy NFG, one also fails to satisfy KFK. Of course, the converse is not true, for it is possible for one to believe truly and fail to know.

Now, and this is the important point here, NFG's account of inference does not only require that the premises one's conclusion depends on be true, but also that they are *justified*. If one were not justified in believing all the propositions on which one's conclusion depends, NFG would not rule out the clearly absurd possibility of someone acquiring inferential knowledge on the basis of mere *lucky guesses*, as when S believes truly each of the premises his conclusion depends on because of individual coin tosses. This possibility is absurd, because knowledge cannot be based on mere lucky guesses. This shows that NFG and KFK are not that different and that, ultimately, how different those accounts are depends on what we take justification to be. If, as it has been suggested,⁴² one is justified in believing that p only if one knows that p, then NFG and KFK would amount to the same thing. If, on the other hand, one can be justified in believing that p even if one

³⁵ Clark (1963).

³⁶ Cf. Goldman (1967: 369–70). Goldman explicitly takes his account of inference to be an improvement on Clark's view on p. 364. There he argues that, if Clark had augmented his view to include beliefs about causal relations as grounds for inferential knowledge, then Clark's view would have been 'almost equivalent' to his view.

³⁷ Cf. Harman (1973, 1980, 1986).

³⁸ Feldman (2002: 33–7).

³⁹ Lycan (2006: 153–8). Lycan also endorsed a similar view back in the 1970s (1977).

⁴⁰ Foley (2012).

⁴¹ To be precise, Clark's version of the view suffers from a problem the other views do not. Since Clark did not distinguish between relevant and irrelevant grounds, his view yields the wrong result in the Harmless Ignorance case (see above).

⁴² Cf. Sutton (2007) and Williamson (2011).

does not know that *p* (like Gettier explicitly assumed to be the case), then NFG and KFK amount to different accounts of inferential knowledge. The point here is that the difference between the two views, although significant, is not irreconcilable.

Lastly, KFK is also entailed by Peter Klein's account of defeasible reasoning. According to Klein, undefeated justification is a necessary condition on inferential knowledge.⁴³ One's belief in the conclusion of one's reasoning is undefeated only if there is no truth *d* such that the conjunction of *d* and the premises in one's reasoning fails to justify one's belief in the conclusion of that reasoning.⁴⁴ This account of defeasible reasoning entails KFK. Suppose, for a *reductio ad absurdum*, that this is not the case; that is, suppose that, for some *S*, *S* knows that *p* as the result of an inference even though she fails to know one of the premises *p* depends on, and that all premises *p* depends on are undefeated. But, if *S* fails to know one of the premises *p* depends on, say *q*, then 'S doesn't know that *q*' is a defeater of *S*'s justification for *p*. This is a contradiction, since we are supposing that there is no defeater of *S*'s justification for *p* and we have uncovered such a defeater. We are then allowed to reject our assumption. It follows that the defeasibility theory entails KFK.

I can also offer a different argument in support of this entailment. If the defeasibility account of *knowledge* is correct, then it seems that there is a genuine defeater of *S*'s justification whenever *S* fails to know one of the premises essentially involved in her inference. The reason for that is straightforward: if *S* is in that situation, then the truth 'S does not know premise *x* of her reasoning' is a defeater of *S*'s justification, because it entails 'Either *x* is false, not believed, or unjustified.' Whatever makes this disjunction true also prevents *S* from knowing the conclusion of her reasoning.

KFK governs inferential knowledge in general. GC is an application of this more general principle to a particular set of cases, the Gettier cases. Accepting KFK was the norm before Gettier. It's time to go back to the good ol' days.

3.1 *Objection: knowledge from non-knowledge*

Before I conclude I will address an issue that suggests that conforming to KFK is not *necessary* for inferential knowledge.

Recently, Ted Warfield,⁴⁵ Peter Klein,⁴⁶ Branden Fitelson,⁴⁷ and Claudio de Almeida⁴⁸ have presented alleged cases of *knowledge from non-knowledge*, cases that allegedly show that one can gain knowledge even though the conclusion of one's inference depends on at least one premise one does not know. Those cases put pressure on KFK as a constraint on epistemically proper reasoning.

I will offer two different types of replies to alleged cases of knowledge from non-knowledge: a non-concessive reply, and a partially concessive reply. As we will see,

⁴³ Cf. Klein (1971, 1981). See also Lehrer & Paxson (1969) and Lehrer (1990).

⁴⁴ The fully worked-out defeasibilist account of *knowledge* has to distinguish between *misleading defeaters* and *genuine defeaters*. See Klein (1979) for the distinction.

⁴⁵ Warfield (2005).

⁴⁶ Klein (2008).

⁴⁷ Fitelson (2010).

⁴⁸ See de Almeida (2003, 2004). De Almeida (2003) was given at the 2003 Meeting of the Central APA.

KFK maintains its status at the center of epistemological theorizing about reasoning and inference, even if we concede some of the points these cases seem to make.⁴⁹

Consider the following alleged case of knowledge from non-knowledge discussed by Warfield.⁵⁰

The Handout Case

Warfield has 100 handouts for his talk; he carefully counts 53 people in attendance and concludes that his 100 handouts are enough. As it turns out he double-counted one person who changed seats during the counting of heads and there are in fact 52 people attending his talk.

This case poses a problem for KFK because Warfield seems to know that 100 handouts are enough even though this knowledge seems to depend on something he fails to know. Let us look first at a non-concessive reaction to this case. This type of reply is non-concessive because it denies that the initial impression that Warfield knows that 100 handouts are enough is probative of him actually knowing that proposition.

The idea is simple enough: Warfield does not satisfy KFK and therefore does not know that 100 handouts are enough. What is more, we can explain why some of us have the intuition that Warfield knows: it is reasonable for Warfield to think he knows there are 53 people in attendance. This gives Warfield an *excuse* to believe this proposition and whatever it entails. But, having an *excuse* to believe truly (or even for *being justified* in believing truly, if having an excuse to believe is sufficient for being justified in believing) is not the same as knowing. So, even though Warfield's belief is highly justified and is, for all practical purposes (such as providing everyone at his talk with a handout), 'as good as knowledge,' he does not know that 100 handouts are enough.

The concept of *excuse* here is a familiar one. One is excused for having ϕ -ed just in case one's epistemic position is such that it is reasonable for one to believe one knows ϕ -ing is permissible. In alleged cases of knowledge from non-knowledge, the protagonist fails to know the conclusion of his reasoning, but it is highly probable, on the protagonist's evidence, that he knows the false premise of his argument. And, since 'I know that p' is highly probable on the protagonist's evidence even though p itself is false, it is *reasonable* for him to believe that p and he has an *excuse* to believe that p. However, it being reasonable for one to believe one's premises are necessary, but obviously not sufficient to generate knowledge of one's conclusion. Thus, the protagonists of cases of alleged knowledge from non-knowledge have inferentially *reasonable* (or justified) belief, rather than knowledge of the conclusion of the argument.⁵¹

⁴⁹ As an aside, I should report that I do not have the intuition that the protagonists in *all* cases discussed by Warfield, Klein, Fitelson, and de Almeida know the conclusion of their reasoning. *Some* of the cases elicit that intuition in me; some *do not*.

⁵⁰ Warfield (2005: 408).

⁵¹ This account of what it is to have an excuse is modeled on Williamson's treatment of excusable but epistemically improper assertions. See Williamson (2000: chapter 11) and Borges (2016).

We are liable to mistake the fact that *it is reasonable for Warfield to believe as he does* for the non-fact that the protagonist in a case of alleged knowledge from non-knowledge knows, because (i) it is highly reasonable for him to believe he knows the premises of his reasoning; (ii) from the protagonist's point of view, there is virtually no difference between knowing that *p* and merely believing truly that *p*, since it is highly likely on his evidence that he knows that *p*; and, (iii) when we assess philosophical cases we tend to put ourselves into the protagonist's shoes and let this influence our assessment of the case.

I can also provide a partially concessive reply to the problem of knowledge from non-knowledge. First, note that, although non-concessive, the first reply did not dispute the idea that Warfield's true belief depends on a false one. The partially concessive strategy disputes this very point. We first concede that there is knowledge in cases of knowledge from non-knowledge, and then insist that, despite appearances to the contrary, this knowledge does not depend on anything the subject is ignorant of. The thought is that alleged cases of knowledge from non-knowledge are disguised cases of harmless ignorance. In Warfield's case this strategy would involve showing that

(t) There are enough handouts

does not depend on the false claim that

(f) There are 53 people in attendance.

Remember that *x* depends on *y* only if *Dependence* is satisfied. It can be plausibly argued that Dependence is not satisfied in Warfield's case. Consider, first, whether *t* *causally* depends on *f*.

(cc) If Warfield had not believed that there are 53 people at his talk, he would not have believed that his 100 handout copies were enough.

The worlds closest to the actual world are worlds in which Warfield's belief about the number, *n*, of people in attendance is such that $n < 100$, and Warfield still forms a belief about the number of attendees. Thus, those are also worlds in which Warfield believes that 100 handouts are enough. Warfield's case also fails the counterfactual test for evidential dependence:

(ec) If Warfield's evidence set did not include 'There are 53 people in attendance,' he would not have been justified in believing '100 handout copies are enough.'

As before, the worlds closest to the actual world are worlds in which Warfield has views about the number, *n*, of people in attendance. In particular, those are worlds in which Warfield thinks $n < 100$. Thus, in the worlds closest to the actual world Warfield's evidence set justifies him in believing that 100 handouts are enough.

How can we know that the worlds closest to the actual world are as I say?⁵² We need two things to evaluate our counterfactuals properly: we need to make the smallest

⁵² Here I am assuming a naive version of the Stalnaker-Lewis semantics for counterfactuals.

number of changes possible to the actual world so that the antecedent of the counterfactual is true; and, second, we need to consider how the world so changed would evolve and whether this evolution would make the consequent true or not. Since the actual world is one in which not only *f* is in Warfield's evidence set, but also one in which Warfield *counts the number of people in attendance*, a world *W* is a member of the set of closest possible worlds (with respect to Warfield's actual world) only if (i), in *W*, *f* is not in Warfield's evidence set and (ii) Warfield counts the number of people in attendance. Such a *W* is closer to the actual world than a world in which we change *both* (i) and (ii). But, since counting the number of people in attendance will lead Warfield to form a belief about the number, *n*, of people in attendance, it is plausible that in the closest possible worlds some proposition about *n* but different from *f* will be in his evidence set and will be enough to justify him in believing that 100 handouts are enough.⁵³

In fact, even if we do not accept the result we arrived at via *Dependence*, there might still be a case to be made that, *in the actual world*, *f* is inessential. Here is why one might think that. First, note that *f*, by itself, is hardly a good enough reason for Warfield to believe that *t*. That is, Warfield's argument for *t* is obviously enthymematic, since it does not even explicitly mention his background (but crucial) knowledge that, for example, he brought 100 handouts to the talk. If Warfield in fact knows that 100 handouts are enough, then this knowledge does not depend exclusively on the proposition that there are 53 people in attendance, for this is not a good enough reason (in and of itself) for anyone in Warfield's situation to believe that 100 handouts are enough. In any plausible way we might fill out the details of Warfield's case, he will have arrived at his talk with evidence relevant to the question of whether the 100 handouts are enough or not. For instance, in any plausible interpretation of the case he arrived at his talk knowing something like the following:

(*t**) If there are between 0 and 100 people in my talk, then 100 handout copies are enough.

If *t** were not part of Warfield's background knowledge (say, because he did not even believe it), it would be hard to explain why he would think *t* followed from *f*, and why he made 100 copies.⁵⁴ The upshot is that, whether we apply *Dependence* or not, Warfield's conclusion does not depend essentially on the falsehood 'there are 53 people at my talk.' The case is, when scrutinized, not a case of knowledge from non-knowledge, but a case of harmless ignorance. Both of these strategies can be extended to cover other cases of alleged knowledge from non-knowledge.

⁵³ Thus, this takes care of the objection, leveled by Luzzi (2014), that without an account of what it is for a proposition to depend on another, one cannot appropriately hold on to KFK when faced with alleged cases of knowledge from non-knowledge. Luzzi also fails to acknowledge the fact that Klein (2008) also offers an account of the relevant notion of dependence.

⁵⁴ One might complain that *t** doesn't partially justify *t* on the ground that the former does not really cause the latter. The problem with this objection is that it overlooks the fact that one's total evidence plays an epistemic role even if not an explicitly causal one.

4 Conclusion

This chapter explored a solution to the Gettier Problem that has been largely ignored by the literature. We argued that this is a mistake and that this solution can be extracted not only from the philosophical tradition, but also from recent accounts of evidence, defeasible reasoning, and from the no-false-grounds account of inference. We also argued that recent cases leveled against the view of inferential knowledge behind the Gettier Conjecture are ultimately unconvincing.

References

- Aristotle. (1994). *Posterior Analytics*, 2nd edition. Oxford: Oxford University Press.
- Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Borges, R. (2016). E=K and the Gettier Problem. *Erkenntnis* doi: 10.1007/s10670-016-9857-z.
- Chisholm, R. (1977). *Theory of Knowledge*. Upper Saddle River, NJ: Prentice-Hall.
- Clark, M. (1963). Knowledge and Grounds: A Comment on Mr. Gettier's Paper. *Analysis* 24(2): 46–8.
- De Almeida, C. (2003). Benign Falsehoods: A Lesson about Inferential Knowledge. Unpublished manuscript.
- De Almeida, C. (2004). *On Useful Falsehoods*. URL: <<http://certaindoubts.com/?p=195>>.
- Descartes, R. (1985). Rules for the Direction of the Mind. In Cottingham, J., Stoothoff, R., & Murdoch, D. (Ed.), *The Philosophical Writings of Descartes*, Volume I. Cambridge: Cambridge University Press.
- Dretske, F. (1971). Conclusive Reasons. *Australasian Journal of Philosophy* 49(1): 1–22.
- Feldman, R. (1974). An Alleged Defect in Gettier Counter-Examples. *Australasian Journal of Philosophy* 52(1): 68–9.
- Feldman, R. (2002). *Epistemology*. New York: Pearson.
- Fitelson, B. (2010). Strengthening the Case for Knowledge from Falsehood. *Analysis* 70(4): 666–9.
- Foley, R. (2012). *When Is True Belief Knowledge?* Princeton, NJ: Princeton University Press.
- Gendler, T. S. & Hawthorne, J. (2005). The Real Guide to Fake Barns. *Philosophical Studies* 124: 331–52.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis* 23(6): 121–3.
- Goldman, A. (1967). A Causal Theory of Knowing. *Journal of Philosophy* 64(12): 357–72.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73(20): 771–91.
- Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.
- Harman, G. (1980). Reasoning and Evidence One Does Not Possess. *Midwest Studies in Philosophy* 5: 163–82.
- Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.
- Hetherington, S. (2011). *How to Know*. New York: Blackwell.
- Jackson, F. (1977). A Causal Theory of Counterfactuals. *Australasian Journal of Philosophy* 55(1): 3–21.

- Kant, I. (1950). *Critique of Pure Reason*. London: Macmillan.
- Klein, P. D. (1971). A Proposed Definition of Propositional Knowledge. *Journal of Philosophy* 68(16): 471–82.
- Klein, P. D. (1979). Misleading ‘Misleading Defeaters’. *Journal of Philosophy* 76(7): 382–6.
- Klein, P. D. (1981). *Certainty: A Refutation of Scepticism*. Minneapolis, MN: University of Minnesota Press.
- Klein, P. D. (2008). Useful False Beliefs. In Smith, Q. (Ed.), *New Essays in Epistemology*. Oxford: Oxford University Press.
- Klein, P. D. (2012). What Makes Knowledge the Most Highly Prized Type of Belief? In Black, T. & Becker, K. (Eds), *The Sensitivity Principle in Epistemology* (pp. 152–69). Cambridge: Cambridge University Press.
- Lehrer, K. (1965). Knowledge, Truth, and Evidence. *Analysis* 25: 168–75.
- Lehrer, K. (1990). *Theory of Knowledge*, 2nd edition. Boulder, CO: Westview Press.
- Lehrer, K. & Paxson, Jr., T. (1969). Knowledge: Undefeated Justified True Belief. *Journal of Philosophy* 66(8): 225–37.
- Lewis, D. (1979). Counterfactual Dependence and Time’s Arrow. *Noûs* 13(4): 455–76.
- Locke, J. (1975). *An Essay concerning Human Understanding*. Oxford: Oxford University Press.
- Luzzi, F. (2014). What Does Knowledge-Yielding Deduction Require of Its Premises? *Episteme* 11(3): 261–75.
- Lycan, W. (1977). Evidence One Does Not Possess. *Australasian Journal of Philosophy* 55(2): 114–26.
- Lycan, W. (2006). The Gettier Problem Problem. In Hetherington, S. C. (Ed.), *Epistemology Futures*. Oxford: Oxford University Press.
- McLaughlin, B. (2006). Is Role-Functionalism Committed to Epiphenomenalism? *Journal of Consciousness Studies* 13: 39–66.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Belknap Press of Harvard University Press.
- Paul, L. A. & Hall, N. (2013). *Causation: A User’s Guide*. Oxford: Oxford University Press.
- Pritchard, D. (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109(3): 247–79.
- Russell, B. (1912). *The Problems of Philosophy*. Oxford: Oxford University Press.
- Russell, B. (1948/2009). *Human Knowledge: Its Scope and Limits*. London: Routledge.
- Scheffler, I. (1965). *Conditions of Knowledge*. Glenview, IL: Scott, Foresman, and Company.
- Shope, R. (1983). *An Analysis of Knowing*. Princeton, NJ: Princeton University Press.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.
- Sosa, E. (2011). *Knowing Full Well*. Princeton, NJ: Princeton University Press.
- Sutton, J. (2007). *Without Justification*. Cambridge, MA: MIT Press.
- Vogel, J. (1990). Are There Counterexamples to the Closure Principle? In Roth, M. D. & Ross, G. (Eds), *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer.
- Warfield, T. A. (2005). Knowledge from Falsehood. *Philosophical Perspectives* 19: 405–16.
- Weatherson, B. (2003). What Good Are Counterexamples? *Philosophical Studies* 115(1): 1–31.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2011). Knowledge First Epistemology. In Bernecker, S. & Pritchard, D. (Eds), *The Routledge Companion to Epistemology* (pp. 208–18). New York: Routledge.