

+

How to Solve the Gettier Problem

James J. Pearce
Adjunct Instructor
Shawnee State University
Portsmouth, Ohio
jpearce@shawnee.edu

The traditional analysis of knowledge, which pre-dates Plato's *Theaetetus*, attempts to state certain conditions which are individually necessary and jointly sufficient for a subject (S) to have knowledge of some proposition (P). Although different philosophers have presented different forms of this analysis, perhaps its most generally recognized form can be stated as follows:

S knows that P	IFF	(i) P is true, (ii) S believes that P is true, and (iii) S is justified in believing that P is true.
----------------	-----	--

Contemporary epistemologists have exhibited an uncharacteristic consensus of opinion in asserting that the traditional analysis, also known as the “justified true belief” theory of knowledge (JTB), cannot be correct. This consensus originates from a paper entitled “Is True Justified Belief Knowledge?” which was published by Edmund L. Gettier in 1963. Gettier's argument, which has had an impact wildly disproportionate to its two-page length, consists of two brief scenarios, each purporting to show an instance of justified true belief that none the less fails to qualify as an instance of genuine knowledge. If Gettier's argument succeeds, then there are at least two instances of justified true belief which are not knowledge. Moreover, in the years since Gettier's paper was published, literally hundreds if not thousands of Gettier-type counterexamples to JTB have been generated by various philosophers. It would seem that the conditions for knowledge which are stated in JTB are *at best* individually necessary but not jointly sufficient for knowledge. If it is possible to state sufficient conditions for knowledge, it would appear that something more is required, something beyond the conditions stated in traditional versions of JTB.

¹ For purposes of clarity, because I will subsequently be discussing a different kind of “epistemic conditions,” I will therefore refer to the sorts of conditions listed here – the traditional conditions of the JTB theory – as “E-conditions.” I will differentiate the other form of conditions when they occur later in the text.

² Note that Gettier problems arise *only* for empirical or “*a posteriori*” arguments which include deductive inferences. To the best of my knowledge, it is not possible to construct a Gettier counterexample for a purely deductive or “*a priori*” argument, because all of the premises in such an argument are true by definition or by fiat, and are therefore *necessarily* true.

The success of Gettier's argument therefore entails the obsolescence of JTB. The "Gettier problem," then, is really only a problem for philosophers who think that Gettier-type counterexamples fail to refute JTB, or that something recognizably similar to JTB can work as an adequate theory of knowledge. As philosophers who support JTB have become a rare species, the furor over the Gettier problem has tended to die down. At least until recently. Since the early 1990s, a new generation of philosophers have begun to appear, among whom the consensus concerning JTB is not so readily apparent.³ Some of these philosophers believe that JTB, or perhaps a version of JTB modified to include additional necessary conditions — what Plantinga calls a "codicil to propitiate Gettier"⁴ — can succeed as an adequate analysis of knowledge. I count myself among this group. In the following paper I will attempt to solve the Gettier problem, but first it is necessary to be clear about exactly what such a solution involves.

1. What does it mean to "solve" the Gettier problem?

Solutions to the Gettier problem can take two forms. First, they can attempt to show that Gettier-type examples fail as counterexamples, and that JTB therefore emerges unscathed. The literature is replete with this kind of counter-counterexample, and such arguments are usually met with counter-counter-counterexamples. Indeed, Shope's survey of the literature discusses over one hundred such examples, most of which were generated in response to contrary counterexamples which had been previously issued.⁵ These arguments and counter-arguments escalate in sophistication and intricacy such that, I believe, a historian could identify discrete waves of argumentation, for example, an infallibilist wave, a defeasibility wave, and so forth. It would seem that, in order to solve the Gettier problem in this first sense, an alleged solution would have to meet every wave of this attack — unless it could be shown that the problem resides in the formulation of Gettier-type counterexamples in general.

³See, for example, Levi, Don S. 1995. "The Gettier Problem and the Parable of the Ten Coins." *Philosophy*. 70: pp. 5-25 [hereafter Levi, 1995]; Bach, Kent. 1998. "Accidental Truth and Would-be Knowledge." <http://userwww.sfsu.edu/~kbach/accidtruth.html>, pp. 1-9 [hereafter Bach, 1998]; Hetherington, Stephen. 1998. "Actually Knowing." *The Philosophical Quarterly*. 48:193. (Located on the internet at http://web4.searchbank.com/infotracs/session/913/486/6898419w5/15!fullart_), pp. 1-14 [hereafter Hetherington, 1998].

⁴ Plantinga, Alvin. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press, p. 182 *et passim* [hereafter Plantinga, 1993].

⁵ Shope, Robert K. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton: Princeton University Press [hereafter Shope, 1983].

The second sense in which the Gettier problem might be solved involves developing a new theory or analysis of knowledge that is not subject to Gettier-type refutation. Not surprisingly, solutions of this second kind also frequently result in the generation of new, higher-level Gettier examples. Since a solution of this type would involve providing a Gettier-proof analysis of knowledge, some philosophers seem to believe that a solution to the Gettier problem would amount to a finished theory of knowledge – the completion of epistemology [Bach, 1998, p.7]. I disagree. I think that the Gettier problem rests on a *misunderstanding of the relation between logical entailment and epistemic justification*. Gettier treats justification as if it were not logically analogous to truth. I shall argue that Gettier, and most subsequent epistemologists, are wrong on this point. If I am correct, the Gettier problem will be dissolved rather than solved, demonstrating that an adequate theory of knowledge is possible. However, demonstrating the *possibility* of such a theory is a far cry from providing the theory itself, which would require, *inter alia*, detailed accounts of justification, evidence, epistemic causality and a resolution of the internalist/externalist debate.

In the remainder of this paper, I shall propose an analysis which I believe functions on two levels – as a dissolution of the Gettier problem and as a proposal for necessary and sufficient conditions for the possibility of knowledge, although these conditions may be rightly criticized as much too general to constitute a completed theory of knowledge. In section (2) I shall analyze and criticize the role typically assigned to valid deductive inference in the generation of classical Gettier examples. In section (3) I will demonstrate how my critique undermines the classical Gettier-type counterexample to JTB. In section (4) I will examine some Gettier-type counterexamples which appear to escape my analysis (non-classical or “new wave” Gettier examples), and argue that they either fail to escape, or they fail to undermine JTB on grounds other than those proposed here. In section (5) I will state my own analysis of Gettier-proof knowledge by adding a condition to the traditional list of necessary and sufficient E-conditions for knowledge.

I may seem to have over-promised here, and so I should be forgiven if I issue two caveats. First, the following analysis applies only to justified true beliefs which are empirical, and which are justified *only* because they are deductively entailed by a set of justified premises. I do not mean to deny that, under normal circumstances, we justify our beliefs by appealing to many things in addition to the manner of their inference; for example, by appealing to evidence, to an internalist ethics of knowing, to an externalist causal connection, etc. But when we are dealing with the Gettier problem, at least in its classical form, we are dealing with anything but normal circumstances, and

therein lies the rub. More on this point later. For the present, suffice it to say that the current analysis will apply *only* to those instances in which it is concluded that, because a subject 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, it therefore follows that S is justified in believing Q.⁶ In other words, justification is subject to deductive syntax. If S adduces further evidence for Q, in addition to the bare deductive inference of Q from P, then the following analysis may but will not necessarily apply to S's claim to know Q.

Second, in the following analysis I will say a great deal about the logical behavior of justification in a deductive inference. However, I will say very little about the nature of justification itself. My reason for circumnavigating the topic of justification is simply that it is difficult to say anything of value about justification in a brief paper not devoted to that topic. For present purposes suffice it to say that, in my view, justification is part of a complex relation (K) which can obtain between a subject or cognitive system (S) and a proposition (P). The relation K may but does not necessarily involve S's possessing evidence for P. Further, K may but does not necessarily involve S's being aware of (or having epistemic access to) K.⁷ Although K may prescribe certain normative conditions for S's knowing P, these normative conditions may, but do not necessarily, impose any moral or quasi-moral duties on S. In effect, I view knowledge (and hence justification) as essentially non-deontological [cf., Plantinga, pp. 6-29]. To paraphrase Plantinga, I will use the term "justification" as a name for that which is necessary for empirical knowledge and (together with truth) sufficient for it *up to* Gettier problems. Moreover, fulfillment of epistemic duty, no matter how fervent and conscientious, is probably irrelevant to knowledge [Plantinga, pp. 27-28].

2. *Deduction and justification.*

⁶This is, of course, Gettier's original description of the relation between the E-conditions of JTB and the knowing subject's acquisition or possession of knowledge; see Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: pp. 121-123 [hereafter Gettier, 1963]. Reprinted in Kleiman and Lewis. 1992. *Philosophy: An Introduction through Literature*. New York, New York: Paragon House, pp. 25-27.

⁷That is, S's knowing P may, but does not necessarily, involve S knowing that S knows P.

In his original article, Gettier makes the following claim: “...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.”⁸ As the history of Gettier examples makes clear, Gettier believes that S is justified in believing Q

regardless of the truth-value of P. So, if P is false (albeit justified), S is still justified in believing Q if Q was deductively (hence validly) derived from P. Why? To my knowledge, Gettier never answers this question, but Lowy attempts to fill certain gaps in Gettier’s notion of justification as follows:

A person is justified in believing a proposition when no more can reasonably be expected of him with respect to finding out whether that proposition is true.... [If] no more can be expected of him in investigating the truth of the matter concerning P, and he deduces Q from P, taking a step which is *truth preserving*, then, no more can be expected of him in investigating the truth of the matter either as far as P or Q is concerned.⁹

This argument is quite odd. A valid deductive argument cannot preserve truth unless the premises are true. Of course, the truth of a premise is not a function of who investigates it, or how much effort is put into the investigation. An uninvestigated claim may be true, and a highly investigated claim may be false. If a premise is false, then a deductive argument based on it cannot preserve its truth, and this fact obtains quite apart from any questions of investigation. *Why, then, should we believe that a truth-preserving (deductively valid) argument preserves justification in circumstances where it does not preserve truth?*

I think the oddness of Lowy’s argument arises from her notion of “reasonable expectation” which suggests a deontological view of knowledge. For example, “What is relevant... is whether or not in collecting and assembling his evidence, he has discharged his responsibility as a truth-seeker...” [Lowy, 1978, p. 107], or similarly, “...if a person is not justified in believing something, then he has failed in some essential duty as a truth-seeker” [Lowy, 1978, p. 107]. Epistemic justification, for Lowy, seems to be like a duty-bound moral action. The results of such moral action may not be what we had hoped, but if the will was good, then the action is moral. Analogously, the results of a belief may be less than knowledge, but if the will was good, then the belief is justified. I maintain, on the contrary, that there is nothing essentially deontological about knowledge. Knowledge is simply a relation between a subject or cognitive system S and a proposition P. If S stands in a certain relation to P, there is knowledge; if S does

⁸Gettier, 1963.

⁹Lowy, Catherine. 1978. “Gettier’s Notion of Justification.” *Mind*. 87: p. 106 [hereafter, Lowy, 1978].

not stand in such a relation to P, there is no knowledge. Knowledge may be useful and hence valuable, but this value is instrumental and not inherently moral, nor is “duty” an essential component of knowing.

To clarify, consider the relation of being “to the left of” something. Suppose a bullet passes through this room. I would find it quite valuable to be “to the left of” that bullet’s trajectory. But this very clear value of being “to the left of” a bullet’s trajectory is not inherent in the trajectory itself, it is rather an accident of my circumstances; likewise with knowledge. Knowledge is not power, except incidentally (*contra* Bacon), nor is it morally laudable, except incidentally (*contra* Plato).¹⁰ In the same vein, justification is not a matter of meeting responsibilities or fulfilling duties. Justification is just one condition necessary in order for S to be related to P in a certain way, and it is possible that this condition may be fulfilled beyond S’s control and external to S’s awareness or epistemic access. To clarify the question, then: what, if anything, warrants the belief that the valid deduction of Q from P is sufficient to justify belief in Q, even if P happens to be a justified *false* belief – or, if Q happens to be true by accident or coincidence, through no direct connection to P?

Validity is the core of first order logic. Proofs and truth-tables are merely ways of showing specific arguments to be valid or invalid. And, of course, what it means for an argument to be deductively valid is that, if all of the premises of the argument are true, then the conclusion must necessarily be true. A conclusion validly derived from true premises cannot possibly be false. This characteristic is usually referred to as the “truth-preserving” property of valid deductive inference. But, in the statement quoted above, Gettier isn’t talking about validity or the deductive preservation of truth. He seems to be talking about something altogether different, something like a “justification-preserving” property of deductive inference. But what is this property? Is it a formal property? If so, how is it defined? Does it entail some necessary connection between truth and justification?

In fact, it is very important for the production of most Gettier-type examples that the concepts of truth and justification be logically independent. That is, it must be possible for some true beliefs to be unjustified, and for some false beliefs to be justified. For this very reason it is not *prima facie* evident that we should assume that the logical behavior of justification is identical to the logical behavior of truth.¹¹ Pending some compelling argument to

¹⁰Knowledge, of course, may be aesthetically valuable, but its aesthetic value is also incidental to its nature.

¹¹In fact, one way of defeating the Gettier problem is to insist that justification entails truth, hence, false beliefs cannot be justified (or “completely justified” or “warranted”) [cf. Sturgeon, 1993, note 14 below]. However, the problem with such a solution is that it leads to infallibilism, i.e., to criteria of knowledge so restrictive that extreme skepticism seems inevitable.

the contrary, therefore, I propose that the only reason to believe that valid deductive inferences are justification-preserving is that, if they aren't, then we don't know much of anything.

However, this bare assertion is hardly a proof that valid deductive inferences are justification-preserving, and we should be wary of appealing to intuition at this point because intuition often resembles self-fulfilling prophecy. It seems necessary, therefore, to impose careful epistemic constraints upon our assumption that valid deductive inference is justification-preserving in arguments *with empirical premises*. I am going to propose what I believe are reasonable constraints of this sort, but first I will attempt to list all possible conditions under which the justification of one or more (empirical) premises might be preserved in deductive inferences to (empirical) conclusions:

- (a) All of the premises are true and justified and the inference is valid.
- (b) All of the premises are true and justified but the inference is invalid.
- (c) All of the premises are true but some premises are unjustified and the inference is valid.
- (d) All of the premises are true but some premises are false and the inference is invalid.
- (e) Some of the premises are false but all of the premises are justified and the inference is valid.
- (f) Some of the premises are false but all of the premises are justified and the inference is invalid.
- (g) Some of the premises are false and some of the premises are unjustified but the inference is valid.
- (h) Some of the premises are false and some of the premises are unjustified and the inference is invalid.

Displayed in the form of a chart or table, all possible combinations of α -conditions would look like this:

Table of α -conditions:

	All claims are true	All claims are justified	All inferences are valid
a)	yes	yes	yes
b)	yes	yes	no
c)	yes	no	yes
d)	yes	no	no
e)	no	yes	yes
f)	no	yes	no
g)	no	no	yes
h)	no	no	no

Since there does not seem to be any formal justification-preserving property which might be attributed to deductive inferences in general, it would be prudent to maintain that deductive inferences are justification-preserving *only* under α -condition (a) above (the first line of the table). There is no obvious reason why justification should be preserved through a deductive inference under any of α -conditions (b)-(h), and I contend that the burden of proof falls to anyone who would maintain the contrary claim.

If it is correct that deductive inferences are justification-preserving only under α -condition (a), then it seems that justification is preserved through a valid deductive inference only when it is piggybacked on truth and validity. When truth is not preserved through the inference, justification is not preserved. So, for example, justification is not preserved under α -condition (b) because the inference is invalid, and truth is therefore not preserved. Justification is not preserved under α -condition (c) because certain essential premises are unjustified, and valid deductive inference cannot generate justification *ex nihilo*. Justification is not preserved under α -condition (d) because some premises are false, so truth cannot be preserved through the inference. And so forth for the remaining α -conditions.

Note that no claim is being made to the effect that false beliefs can never be justified. It is certainly possible that false beliefs may be justified, but the claim here is that the justification of such beliefs cannot be preserved *through deductive inference* when truth is not preserved. Hence, if P is false but justified, and Q is deduced from P, then the justification of P cannot be transferred through the inference to Q. Why? Because the preservation of justification through a deductive inference is contingent upon the preservation of truth through that inference. Where there is no truth, truth cannot be preserved, and for the same reason, neither can justification. If Q is deduced from a false premise, then the justification of Q will require independent evidence, that is, *evidence which justifies Q independently of any reference to the justification of P*.

3. Countering counterexamples.

It seems appropriate to begin with Gettier's own arguments. In **Gettier's first example**, he says that Smith starts from the following premise:

(i) Jones is the man who will get the job, and Jones has ten coins in his pocket.

Gettier points out that (i) entails (j):

(j) The man who will get the job has ten coins in his pocket.

Proposition (j) is of course true, although proposition (i) is false because Jones doesn't get the job; the job goes to someone else with ten coins in his pocket. Note that Smith has *absolutely no independent evidence for (j)*; that is, over and above his evidence for (i) – which consisted in part of counting the coins in Jones's pocket – Smith has no evidence for (j). But, according to Gettier, Smith is justified in believing (j). Why is Smith justified in believing (j)?

¹² I will refer to these conditions as " α -conditions" in order to distinguish them from the previously described "E-conditions."

Because (j) is validly deduced from (i), and (i) is justified. However, it has been argued here that the inference from (i) to (j) is not justification-preserving unless (i) is true, justified, and the inference from (i) to (j) is valid. Since (i) is not true, the justification of (i) does not flow through the inference, so (j) is not justified. Consequently, (j) is not an instance of justified true belief, and Gettier's Case I fails as a counterexample to JTB. **Note that Case I is an attempt to preserve justification under α -condition (d) above.**

In **Gettier's second example** we again find the epistemically hapless Smith, who is now justified in believing the following proposition:

(k) Jones owns a Ford.

By a valid application of the rule of addition, Smith constructs the following proposition (bear in mind that Smith knows nothing of Brown's whereabouts):

(l) Either Jones owns a Ford or Brown is in Barcelona.

As it turns out, (k) is false and (l) is fortuitously true because, by some wild coincidence, Brown actually is in Barcelona! Gettier clearly believes that (l) is justified because (k) was justified, and (l) was validly deduced from (k). But is (l) really justified?¹³ Not if the justification-preserving properties of deductive inference are restricted to α -condition (a). If (k) is false, then the justification of (k) is not preserved in the inference to (l), so (l) is not an example of justified true belief. Gettier's Case II (which again tries to preserve justification under α -condition (d)), also fails to provide a counterexample to JTB.

The Gettier-Lehrer problem has received almost as much press as the original version, so we should see how it fares.¹⁴ In this scenario, I am completely justified in believing:

(m) Nogot owns a Ford.

Based on this justified belief, and the fact that I share an office with Nogot, I infer:

(n) Someone in this office owns a Ford.

As it turns out, Nogot owns no Ford, but Havit (who happens to be in the same office as Nogot and myself) does indeed own a Ford. It seems, therefore, that (n) is true because Havit is present in the office and owns a Ford, and

¹³ Note again that the only justification for (l) is the valid deductive inference from (k). Aside from this inferred justification, there is no independent evidence to support (l).

¹⁴ Lehrer, Keith. 1974. *Knowledge*. Oxford: Clarendon Press, pp. 219-220 [hereafter, Lehrer, 1974].

(n) is justified because (n) is validly inferred from (m).¹⁵ But if the deductive preservation of justification is confined to α -condition (a), then the justification of (m) does not survive the inference to (n). Once again (n) fails as an example of justified true belief and hence as a counterexample to JTB. (This example, too, tries to preserve justification under α -condition (d).)¹⁶

Turning to a more troublesome, new-wave example, Zagzebski argues as follows: Dr. Jones, a physician, has very good inductive evidence that her patient, Smith, is suffering from virus X. These symptoms are not compatible with any other known virus, all of the evidence upon which Jones bases her diagnosis is accurate, and there is no evidence accessible to her which counts significantly against this conclusion. None the less, Smith's symptoms are due to a distinct and unknown virus Y. What makes this a Gettier-type example, however, is that unlucky Smith has just contracted virus X, but has contracted it so recently that no symptoms are yet exhibited.¹⁷ Zagzebski's argument can be schematized as follows:

- (o) Smith's symptoms are not compatible with any known virus except X,
- (p) All of the available evidence is true, and
- (q) There is no contrary evidence available, therefore
- (r) Smith has contracted X.

Note that the conjunction of (o), (p) and (q) does not entail (r). Therefore, this argument is either an enthymeme and one or more premises have been suppressed, or the inference to (r) is truly non-deductive.¹⁸ If the inference is non-deductive, then the current analysis does not apply to it. Conversely, if the argument is an enthymeme, and if the missing premise is false (if, for example, it is a premise like "There is no unknown virus which could cause these

¹⁵ By now the reader has no doubt become adept at noticing that, except for the justification inferred from (m), there is no independent justification of (n). If the justification of (n) is not legitimately inferred, then (n) is not justified.

¹⁶The problem, therefore, is not that Existential Generalization is an invalid or illegitimate argument form; the problem is, rather, that *deductive validity is not sufficient in an empirical argument to preserve justification unless it also preserves truth*. For more on the debate over the epistemic legitimacy of Existential Generalization, see New, C.G. 1966, "Some Implications of 'Someone'," *Analysis* 26: pp. 62-64, and Smith, James M. 1966, "New Implications of 'Someone', II," *Analysis* 26: pp. 207-208.

¹⁷ Zagzebski, Linda. 1994. "The Inescapability of Gettier Problems." *The Philosophical Quarterly* 44:174, pp. 71-72 [hereafter Zagzebski, 1994].

¹⁸ A non-deductive inference could be inductive or abductive, i.e., an inference to the best explanation, cf., e.g., Douven, Igor, "Abduction", *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/abduction/>>.

symptoms”), then justification can be preserved in the conclusion only under α -condition (d), a claim which has already been disputed.

Suppose, however, that all of the premises are true and justified, and that the premises entail the conclusion. It seems that, under these circumstances, Dr. Jones should *know* that Smith has contracted X, but Zagzebski seems to think not. The reason Zagzebski believes this is that the conclusion is true by accident, hence the truth of the conclusion is somehow unrelated to the justification of the premises. If Dr. Jones “knows” that Smith has contracted X, it is by mere coincidence, pure luck. But can Dr. Jones *know* that Smith has contracted X if Jones’s diagnosis is right by chance? Many epistemologists argue that such a conclusion would be “...unacceptable. Luck does not contribute to our possession of knowledge.”¹⁹ If Zagzebski’s example is construed in this way, and if it is in fact a genuine Gettier-example (that is, if Dr. Jones really does not know that Smith has contracted X), then it seems to be a type of Gettier-example which escapes the present line of analysis. Let us pause to examine this “new wave” kind of Gettier-example.

4. *Gettier examples that escape this analysis.*

As each new wave of Gettier-type counterexamples became progressively more and more sophisticated, it was inevitable that philosophers would notice the role played by falsity in undermining Gettier-type arguments against JTB, so it was inevitable that someone would try to come up with Gettier examples which contain no false premise or belief. If such examples are indeed possible, then they would clearly escape the force of the present argument. In this section, however, I will consider two such examples in an attempt to demonstrate that even if they sidestep the present analysis, it is only because they fail on grounds entirely different from those so far discussed, and hence are in fact genuine cases of S’s knowing that P.

The first example was produced by Richard Feldman, and it is a variation on Lehrer’s theme.²⁰ It is given that Nogot owns no Ford, Havit does, and Nogot, Havit and Smith occupy the same office. Suppose, then, that Nogot tells Smith that she owns a Ford and even shows her a title to that effect. Suppose, further, that up till now Nogot has always been reliable and honest in her dealings with Smith. Typically, at this juncture, Smith would

¹⁹ Sturgeon, Scott. 1993. “The Gettier Problem.” *Analysis* 53.3: p. 161 [hereafter Sturgeon, 1993].

²⁰ Feldman, Richard. 1974. “An Alleged Defect in Gettier Counter-examples.” *Australasian Journal of Philosophy* 52:1, pp. 68-69 [hereafter Feldman, 1974].

generalize her false belief that Nogot owns a Ford and infer that someone in the office owns a Ford, a maneuver which would fail under the exclusion of α -condition (d) since Nogot is lying. However, Feldman's strategy is significantly different,²¹ and the inference looks like this:

- (s) Nogot presents evidence that she owns a Ford,
- (t) Someone in the office presents evidence that she owns a Ford, hence
- (u) Someone in the office owns a Ford.

It follows that (s) is a true statement (the fact that Nogot's evidence is specious is irrelevant), (t) is a true statement, and (u) is a true statement. Moreover, the inference from the conjunction of (s) and (t) to (u) is justified by Nogot's prior credibility. Nevertheless, Feldman concludes, Smith does not know (u). Such a counterexample, if it works, introduces no element of falsity into the inference. But does it work?

As in Zagzebski's example, the conjunction of (s) and (t) clearly does not entail (u) – presenting evidence for β does not entail that β is true. Once again, whatever kind of inference this is, it is either enthymemic or non-deductive.²² The analysis of this example proceeds parallel to that of Zagzebski's: First, if it is truly non-deductive, it is beyond the scope of this paper. If it is enthymemic and the suppressed premise is false, then Feldman's example is a case of justification under α -condition (d), which is being disputed.²³ Second, if the argument is enthymemic and all of the implicit and explicit premises are true and justified, and the premises entail the conclusion, then does it follow that Smith fails to know (u)? And if so, *why* does it follow? We will see how philosophers attempt to answer this question, but first one more example.

The next example is taken (slightly modified) from Gilbert Harman.²⁴ In Harman's scenario, S believes a bylined article, written by an eyewitness, read in a reliable paper, which reports that a famous civil-rights leader has

²¹I am compressing Feldman's argument, but I do not believe it affects the argument's cogency.

²²If this inference is somehow supposed to be deductive as it stands, then it is invalid. To maintain that justification survives an invalid deductive inference (with premises which are true and justified) would be to assert that deduction under α -condition (b) is justification-preserving, a claim which has already been disputed.

²³In this particular instance, it seems at least likely that Feldman's example really is an enthymeme with one or more false suppressed premises. The reason is that Smith's evidence consists of false claims made by Nogot. Surely these false claims must figure somewhere into Smith's reasoning, or perhaps a claim like "I can trust good ol' Nogot, she's always been reliable." This latter claim is of course false. Nogot may have been trustworthy on prior occasions, but the fact is that this has changed.

²⁴Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press, pp. 143-144 [hereafter Harman, 1973].

been murdered. Unknown to S, many people are skeptical about this report because later reports (which S has not seen) have denied that the assassination occurred. However, the later reports were the product of a secret conspiracy intended to prevent racial violence and rioting. If S had seen the false contrary evidence, it would have undermined S's justification for believing, and hence S would not have *known* that the civil-rights leader was dead. It was simply a stroke of luck that S encountered none of this evidence. However, if S's knowing that P is a matter of brute luck, does S know P?

Examples like Feldman's and Harman's introduced into the Gettier debate a wave of **counterfactual analysis**. Such analyses proceed as follows: There is a possible world (W1) in nearby logical space. In W1, S acquires the false evidence (for example, S reads the false reports denying the assassination, or Smith foolishly takes Nogot at her word). S consequently lacks knowledge in W1. But if W1 is sufficiently close in logical space to the actual world (W), then S lacks knowledge in W. The closer W1 is to W, the more it looks as if S's knowledge is dependent on luck. Although this kind of analysis has been extremely popular in some recent literature, it is fraught with difficulties and subject to convincing criticism offered in an article by **Stephen Hetherington** [Hetherington, 1998]. According to Hetherington, to deny that S knows P in W because S does not know P in W1 is to confuse lacking knowledge with almost lacking it [Hetherington, 1998, p. 8]. Such confusion results from the supposition that the counterfactual lack of knowledge implies the actual lack of it [Hetherington, 1998, p. 11], and Hetherington argues that this kind of analysis is guilty of what he calls the "epistemic counterfactuals fallacy" (ECF) [Hetherington, 1998, pp. 3, 8ff].²⁵

If Hetherington is right, one commits ECF in maintaining that S does not know P because S might not have known P under different circumstances (or in a different possible world). Such claims are like saying that I cannot know that the object before me now is a desk, because it might have been a duck. S's knowing P, therefore, seems to be dependent on S's actual situation or circumstances. But if knowledge is circumstantial or situation-dependent in this way, then the conclusion that knowledge involves luck seems inescapable.²⁶ Indeed, Hetherington maintains that

²⁵A belief which is counterfactually false counterfactually fails to be knowledge, but to assume that a counterfactual lack of knowledge implies an actual lack of knowledge is to commit ECF [Hetherington, 1998, p. 8].

²⁶The term "luck" remains vague in all of these discussions. Participants in the debate seem to assume that it is possible to define "luck" in some meaningful way, which in turn assumes on another level that we do not inhabit a deterministic or fatalistic universe.

denying luck a role in knowledge amounts to asserting an infallibilist principle [Hetherington, 1998, p. 9], and that fallibilism is necessarily compatible with knowledge involving luck [Hetherington, 1998, p. 12]. “There can be close possible worlds where a person’s knowledge that P disappears... without its therefore disappearing from this world. But a fallibilist should allow that knowledge can be fragile, or contingent, or lucky...” [Hetherington, 1998, p. 13]. S can know P even if S might not have known P under different circumstances; in short, S can luckily know P [Hetherington, 1998, p. 8]. Perhaps it is just luck that the object before me now is a desk and not a duck. The kind of Gettier examples presented by Harman and others, therefore, fail to defeat the preceding analysis because they are not really Gettier examples at all, at least, not in the form that Gettier presented. Hence, the possibility that S might not have known P under different circumstances or in a different possible world presents no evidence against S’s claim to know P here and now. So, as a matter of fact, S really does know P, at least defeasibly.

5. *A fourth condition for knowledge.*

In section (1) I issued a promissory note stating that I would offer a fourth E-condition which, in conjunction with the traditional three E-conditions, would be sufficient for Gettier-proof knowledge. The time has come to make good on this promise. The following are the E-conditions which I believe are individually necessary and jointly sufficient for knowledge.

- S knows that P IFF (i) P is true,
- (ii) S believes that P is true,
- (iii) S is justified in believing that P is true, and
- (iv) If P is justified through deductive entailment alone, then α -condition (a) is satisfied; otherwise, deductive entailment is not sufficient to justify P.

Note that E-condition (iv) differentiates between methods of justification for P. The thrust of (iv) is that if P is justified *solely* by its deductive derivation from another justified proposition, and if P is an empirical claim, then P must be deduced exclusively through α -condition (a)-type inferences; otherwise, P is not justified. If some of the inferences involved in the entailment of P are not, or are not known to be, α -condition (a)-type inferences, then, either P is not justified or P must be justified in some way independent of this deductive entailment. Or, again, if there is no independent evidence for P, and if P is alleged to be justified solely on the basis of being entailed by some

other justified statement (R), then E-condition (iv) tells us that the inference from R to P is sufficient to justify P if and only if R is true, R is justified, and the inference itself is valid. Otherwise, deductive inference alone, *without further independent evidence*, is not sufficient to guarantee the justification of P.

It is impossible to anticipate all the objections to this analysis of knowledge, but two problems come immediately to mind: 1) It is meaningless to insist that P be justified independently of all logical entailments, because any justification of P will involve inferring P from some claim or other. While this point is certainly correct, some inferences leading to P may be non-deductive. Moreover, all instances of the classical Gettier problem are constructed in such a way that a true (empirical) conclusion is inferred from a determinate set of justified false (empirical) claims. This restriction of the set of claims from which P can be inferred is patently artificial.²⁷ If E-condition (iv) is taken to mean that any empirical P should be justified independently of deductive entailments, then this claim should be interpreted as, “...independently of any *fixed or determinate set* of deductive entailments.” In other words, it should be possible to infer P from various, logically discrete sets of justified empirical claims, some of which are likely to be true. *Put as a rule of thumb: When reasoning through a series of logical steps involving empirical premises, make sure each step is separately justified. Don't count on the justification of previous steps to carry over into subsequent steps.* I do not believe that this constraint is alien to the way we go about justifying knowledge-claims in the real world. In fact, I think that – especially in the empirical sciences – we routinely come at our conclusions from different directions before claiming to know them. Hence, Smith needs to seek further evidence and re-examine his conclusions after having drawn them; had he done so, it would have become obvious that his beliefs were unjustified.

2) It is not usually possible to know whether P has been inferred using only α -condition (a)-type inferences. The reason for this difficulty is that condition (a) requires that the premises from which P is inferred are not only justified, but also *true*, and it can be notoriously difficult to know whether a claim or belief is true.²⁸ (That, after all, is why we have a “problem of knowledge” in the first place.) Here I grant the full force of the objection, but I do not

²⁷ If I were following Gettier's notation strictly, I should be using the variable 'Q' at this point instead of 'P'; however, I was afraid that switching suddenly from talk of 'P' to talk of 'Q' would confuse the reader, so I have chosen to remain consistent in my own terminology instead of Gettier's.

²⁸ To reiterate, this difficulty does not extend to the *a priori* or “exact” sciences where axioms may be true by stipulation, fiat, or definition; hence, I do not believe it is possible to construct a Gettier-counterexample for a deductive argument in which all of the premises are true *a priori*, nor do I believe that the present analysis of knowledge creates any undue problems for the possibility of knowledge within the ‘pure’ or ‘exact’ disciplines.

see it as an objection against this analysis of knowledge. Rather, I see it as *an epistemic caution against using deductive inference from a determinate set of justified premises as the sole basis for justifying an empirical knowledge-claim*. Bach says that “...what, beyond being justified, makes a belief warranted is that it not be in circumstances that make it liable to being a Gettier example” [Bach, 1998, p.7]. E-condition (iv) enables us, I believe, to specify fairly simply what it is that makes a belief “liable to being a Gettier example,” to wit, *over-reliance on deductive inference as a source of empirical justification*.

I have argued that the analysis of knowledge given above blocks any Gettier problem constructed along classical lines, since these problems rely heavily on the deductive transference of justification from false empirical statements to inferred true empirical statements. I have also argued, in section (4) above, that “new wave” counterfactual Gettier problems, such as those proposed by Zagzebski, Feldman, and Harman, are not true Gettier problems at all.

If the arguments presented here go through, then the Gettier problem has been solved, or perhaps dissolved. But what exactly is the cash value of this dissolution? Clearly the present analysis has contributed only minimally to the theory of justification, no new theory of evidence has been advanced, and no contribution to the internalist/externalist debate has been made. We seem to be no nearer than ever to the “completion” of epistemology.

Ultimately, I think that the dissolution of the Gettier problem teaches three valuable lessons: 1) Knowledge is possible and probably does exist. 2) It is fallacious to differentiate the logical behavior of justification in a deductive inference from the logical behavior of truth in a deductive inference. Justification is preserved through a deductive inference if and only if truth is preserved through the inference, and not arbitrarily inserted into the inference through luck or coincidence. In those circumstances where truth is not preserved – either because one or more premises are false, or because the inference itself is invalid – then justification is likewise not preserved. 3) For precisely this reason, deductive inference alone is probably not the best possible choice as sole grounds for the justification of belief.

Works Cited

Bach, Kent. 1998. "Accidental Truth and Would-be Knowledge." (Located on the internet at <http://userwww.sfsu.edu/~kbach/accidtruth.html>).

Douven, Igor, "Abduction", The Stanford Encyclopedia of Philosophy. (Summer 2017 Edition)

Feldman, Richard. 1974. "An Alleged Defect in Gettier Counter-examples." *Australasian Journal of Philosophy* 52:1.

Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23.

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

Hetherington, Stephen. 1998. "Actually Knowing." *The Philosophical Quarterly*. 48:193. (Located on the internet at <http://web4.searchbank.com/infotrac/session/913/486/6898419w5/15!fullart>).

Lehrer, Keith. 1974. *Knowledge*. Oxford: Clarendon Press.

Levi, Don S. 1995. "The Gettier Problem and the Parable of the Ten Coins." *Philosophy*. 70.

Lowy, Catherine. 1978. "Gettier's Notion of Justification." *Mind*. 87.

New, C.G. 1966, "Some Implications of 'Someone'," *Analysis* 26.

Plantinga, Alvin. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press.

Shope, Robert K. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton: Princeton University Press.

Smith, James M. 1966, "New Implications of 'Someone', II," *Analysis* 26.

Sturgeon, Scott. 1993. "The Gettier Problem." *Analysis* 53.3.

Zagzebski, Linda. 1994. "The Inescapability of Gettier Problems." *The Philosophical Quarterly* 44:174.