Australasian Journal of Philosophy

ISSN: 0004-8402 (Print) 1471-6828 (Online) Journal homepage: https://www.tandfonline.com/loi/rajp20

Defeasibility and Gettierization: A Reminder

Claudio de Almeida & J.R. Fett

To cite this article: Claudio de Almeida & J.R. Fett (2016) Defeasibility and Gettierization: A Reminder, Australasian Journal of Philosophy, 94:1, 152-169, DOI: 10.1080/00048402.2015.1009127

To link to this article: https://doi.org/10.1080/00048402.2015.1009127

Australasian Journal of Philosophy, 2016
Vol. 94, No. 1, 152–169, http://dx.doi.org/10.1080/00048402.2015.1009127

Routledge
Taylor & Francis Group

DEFEASIBILITY AND GETTIERIZATION: A REMINDER

Claudio de Almeida and J.R. Fett

For some of us, the defeasibility theory of knowledge remains the most plausible approach to the Gettier Problem. Epistemological fashion and faded memories notwithstanding, persuasive objections to the theory are very hard to find. The most impressive of those objections to the theory that have hitherto gone unanswered are examined and rejected here. These are objections put forward by Richard Feldman, Richard Foley, and John Turri. While these are all interesting, the objection recently put forward by Turri is, we think, by far, the most serious threat to the theory that we have seen in a long time. A successful reply to it requires a surprising amount of care, as it turns out. If tenable, Turri's objection deals a devastating blow to the theory developed by Roderick Chisholm, Keith Lehrer, Peter Klein, Marshall Swain, Risto Hilpinen, John Pollock, and Paul Moser, among others. Under scrutiny, however, the threat proves illusory. It results from inattention to a crucial, but relatively subtle, aspect of the theory. Interestingly, there is only one source in the defeasibilist literature for a precise account of this crucial feature of the theory: one of the most neglected passages in Peter Klein's work on the issue. That crucial feature is put under the spotlight here. Our response to three major objections to the defeasibility theory requires a brief introduction to the theory as an anti-Gettier weapon, an introduction aimed at countering the numbing simplicity that characterizes most introductions to the topic. Following this brief introduction, those three objections are tackled. We conclude, on that basis, that anybody who fails to notice how resilient the defeasibility theory has proven to be for the last fifty years has a seriously deficient understanding of the current state of play in the debate over the Gettier Problem.

Keywords: Gettier problem, defeasibility theory of knowledge, knowledge, epistemic justification, Peter D. Klein

1. Introduction

Not too long ago, Jonathan Kvanvig [2005] claimed that, among all the proposals aimed at solving the Gettier Problem that make room for a notion of epistemic justification, the defeasibility theory of knowledge stands out as 'the only interesting game in town'. We agree. Some of our most influential epistemologists disagree. These include Richard Feldman [2003], Richard Foley [2012], and John Turri [2012]. If they remain unanswered, the objections to the theory put forward by these philosophers will seem to show that the theory developed by Keith Lehrer, Peter Klein, Marshall Swain, Risto

Hilpinen, and John Pollock, among others, is a complete waste of time, no less.[1] These are objections designed to hit the theory like a clean knockdown punch—that is, without incurring the exorbitant costs paid by some of the other foes of the theory, such as embracing infallibilism, or explaining knowledge without reference to a justification condition, or decrying the traditional focus on belief states, or any of the other dramatic moves with which we have become familiar in Gettier's wake.[2] For the defeasibility theory, there is no surviving these objections, if they work. In what follows, we show you that they don't.

While they are all *prima facie* impressive, some of these objections are harder to respond to than others are. We submit that Turri's is the hardest of the bunch. He thinks that, on close inspection, the theory proves to be weaker than anybody had hitherto imagined. According to him, when properly understood, the theory commits us to the view that some paradigmatic Gettier cases turn out to be cases of knowledge. (In fact, Turri's objection commits us to the view that *every* Gettier case is a case of knowledge, as we shall shortly see.) But that's clearly absurd. *Ergo*, the theory is false. Turri's *reductio* fails, however, as we show you in section 5 below. But we think that the failure is highly instructive, as it calls attention to a crucial, but relatively subtle, aspect of the theory. In order to put the error in sharp relief, we turn to the work of Peter Klein. It is generally agreed that no other author has done as much for the defeasibility theory as Klein has done.[3] And it was Turri's failure to appreciate that key aspect of Klein's work on the theory that accounts for the failed refutation. In order to show you how Turri's—or *anybody's*—perception of the theory is crippled by the oversight, we provide you with a brief introduction to the theory as an anti-Gettier weapon. Our introduction (sections 2 and 3)—designed, as it is, to highlight elements of the theory that have been largely neglected, or misunderstood, elsewhere—paves the way for the removal of the objections in sections 4 and 5 below.

2. Gettierization: Two Cases for the Road

For the discussion that follows, we should bear the essentials of Gettierization clearly in mind. Gettier cases are a varied lot. But two of the clearest Gettier cases to be found in the literature on the problem should suffice for our purposes here. Consider, to begin with, the following instance of Gettier's original recipe, Turri's [2012: 215] variation on Lehrer's famous [1965] 'Mr. Nogot' case:

> LAMB: One of Dr. Lamb's students, Linus, tells her that he owns a Lamborghini. Linus has the title in hand. Dr. Lamb saw Linus arrive on campus in the Lamborghini each day this week. Linus even gave Dr. Lamb the keys and let

[1] Some of the key references are Lehrer [1965], Lehrer and Paxson [1969], Klein [1971, 1981, 1996, 2008], Swain [1974, 1981], Pollock [1986], and Hilpinen [1988].
[2] For reviews of the Gettier literature, see Robert Shope [2002, 2004], and Turri [2012].
[3] For evidence of Klein's leadership among defeasibility theorists, see Plantinga [1996], Swain [1996], and Shope [2002, 2004]. And notice the very important developments in Klein [2008], where his leading role in this research project is reasserted.

> her take it for a drive. Dr. Lamb believes that Linus owns a Lamborghini, and as a result concludes, 'At least one of my students owns a Lamborghini'. As it turns out, Linus doesn't own a Lamborghini. He's borrowing it from his cousin, who happens to have the same name and [birthdate]. Dr. Lamb has no evidence of this deception, though. And yet it's still true that at least one of her students owns a Lamborghini: a modest young woman who sits in the back row owns one. She doesn't like to boast, though, so she doesn't call attention to the fact that she owns a Lamborghini.

A very conservative view of cases such as LAMB would have us believe that what explains the fact that Dr. Lamb does not know that one of her students is a Lamborghini owner is the fact that a false belief—the belief that Linus owns the car—plays an essential role in Lamb's reasoning. As Turri rightly notes, one does not need a defeasibility theory in order to explain how knowledge fails to result from reasoning in which a false belief plays an essential role. One need only appeal to the Aristotelian assumption (not identified as such by Turri) according to which only knowledge yields knowledge in inference. That assumption has until very recently been assimilated by every response to the Gettier Problem;[4] which leads us to the familiar no-false-lemmas view of Gettier cases.

But the no-false-lemmas view fails to account for cases where one's premises are all cases of knowledge and where the reasoning is formally unimpeachable. Such cases are relatively hard to find in the Gettier literature.[5] Turri offers us the following, very useful, instance of that Gettier recipe. Here is our adaptation of it:

> LUCKY LAMB: One of Dr. Lamb's students, Linus, tells her that he owns a Lamborghini, and offers Dr. Lamb abundant evidence that he owns the car. (He shows her the title, gives her a ride, parks the car in his garage, etc.) Dr. Lamb is convinced: Linus does own a Lamborghini, she now believes. What she doesn't know about Linus is that he has been insincere. He promised to a cousin of his who happens to have the same name and birthdate that he would take care of the cousin's car while the cousin is on a trip. He then concocted the story about being a Lamborghini owner in order to impress Dr. Lamb. But, unbeknownst to Linus, his cousin died shortly afterwards, and the cousin's unopened will gives Linus the car. Although he did not know it when he put on his show of wealth, Linus has owned the Lamborghini ever since he laid hands on it.

We will expect you to agree with Turri and us that, in LUCKY LAMB, Dr. Lamb remains ignorant while having the true and justified belief that Linus owns a Lamborghini. With conniving, insincere, Linus as her source, she believes what he himself deems false. But luck intervenes, and Gettierization is consummated.

Interestingly, in LUCKY LAMB, there is no need to assume that a false belief is playing a role in Dr. Lamb's reasoning to the conclusion that Linus is a Lamborghini owner. One may be tempted to pollute the case with any

[4] We now have reason to believe that the Aristotelian thesis is false. See Klein [2008] and de Almeida [forthcoming]. The complication should not occupy us here, as it gives rise to an objection that has been tackled elsewhere, and we are focusing only on *unanswered* objections.

[5] A Gettier case involving inference only from true beliefs was originally offered in Feldman [1974].

number of false beliefs about the deception. For instance, one may see the case as one where she uses the belief that Linus is a sincere testifier as a premise. But nothing of the sort is integral to the case. She need have no such false belief. There is a clear inferential path from a set of beliefs, all of which are true, about Linus's Lamborghini-owning behaviour to the Gettierized conclusion.[6] So, Turri is entirely right: the no-false-lemmas view is too weak to work here. Enter the defeasibility theory.

3. Defeasibility and Gettierization: A Crash Course

The defeasibility theory evolved over time in response to the pressure from Gettier-type cases, from its humble, Gettier-oblivious, beginnings in the work of Roderick Chisholm [1964] to its most sophisticated incarnation in Klein [1981]. As an anti-Gettier weapon, its conceptual core is as simple as it is appealing: justification that is good enough for knowledge withstands the addition of true beliefs to the agent's doxastic system.[7] If the truth destroys your justification for a given belief, there is something deeply wrong with that justification. According to the theory, that is exactly what Gettier cases show.[8] In those cases, try adding relevant truths to the believer's doxastic system and you will soon notice that she is no longer justified in holding that true belief that is not a case of knowledge. When that happens, we call the justification-busting truth a 'defeater' of that justification.[9]

Defeaters are selected through a little thought-experiment. First, you think of a Gettier case, for instance, LAMB. There, you find a justified true belief that is not knowledge, namely 'At least one of my students owns a Lamborghini.' (As usual, let's call this the 'target belief'.) Second, you easily notice that there is a relevant truth not believed by the victim, Dr. Lamb, which is such that, if added to Lamb's belief system, destroys—becomes an *overrider* of—her justification for the target belief. Most obviously in the case, that truth is (expressed by) 'Linus does not own a Lamborghini.' This bears out the impression that the victim is being misled by the evidence that she has for the target belief. Surely, anybody reasoning from that misleading evidence is justified in having the target belief. Dr. Lamb's reasoning is unimpeachable. It's the misleading environment created by Linus that robs her from knowing that she has a Lamborghini owner among her students. If you conjoin the defeater (turned-overrider) with Lamb's evidence, the resulting conjunction most definitely does not justify the target belief. Lamb's justification is not *truth-resistant*, in the sense of being unable to withstand the addition of truths to her belief system. The defeasibility theory efficiently

[6] Lamb might abductively reason as follows: 'Linus behaves just like a Lamborghini owner, and I have no evidence that he's a deceiver; therefore, he owns the Lamborghini he drives.'

[7] In contemporary epistemology, we are alerted to this fundamental thought when Jaakko Hintikka [1962: 18] notes that one who (sincerely) claims to know a given fact 'implicitly denies that any further information would have led him to alter his view'—by which he clearly means that knowledge must be compatible with more knowledge.

[8] A crucial caveat, however, will be added to this claim in section 5 below, in our reply to Turri's objection.

[9] Defeaters are, by definition, true propositions not in the agent's 'belief box'. Beliefs that are counterevidence to a given belief are labelled, by Klein, as 'overriders' of the justification of that given belief. Unfortunately, however, many have become used to calling such beliefs 'internal defeaters', or 'psychological defeaters', or even simply 'defeaters'! That's the clumsy but popular way with the terminology.

explains why a justification that is irreproachable from the internal point of view may not be truth-resistant in a misleading environment. The fundamental assumption is that truth-resistance is a necessary condition on knowledge-yielding justification. Unlike the justified belief that is Gettierized, true belief arising from truth-resistant justification is not incompatible with *indefinitely many more true beliefs*.[10] In LAMB, the justification for the target belief is incompatible with true beliefs about Linus's deception, since these true beliefs obviously justify disbelieving that Linus owns the car. There, you can have more true beliefs at the cost of not having only justified beliefs. Knowledge-yielding—that is, truth-resistant—justification welcomes true beliefs without end. But Gettierized justification is selective: you can only have so many true beliefs on a given occasion without some of them becoming unjustified.[11]

The theory came of age when Keith Lehrer and Thomas Paxson [1969] noticed that the defeasibility theory that accounts for Gettierization in cases such as LAMB is too strong. Clear cases of knowledge would, on that early version of the theory, turn out to be cases of Gettierized belief. Here's our version of their famous case (a case we shall refer to in what follows).

> DEMENTED: You see a man who looks just like Tom Grabit stealing a book at the library. Your observation of the man and his demeanour (under good lighting, at short distance, and drawing on good memory of what Tom looks like) leaves no room for reasonable doubt: (P) the thief must be Tom Grabit. Unbeknownst to you, however, (D) Mrs. Grabit, Tom's mother, claims that, while Tom is away on a trip, an identical twin of his is at the library on the day in question. But, still unbeknownst to you, (R) Mrs. Grabit is an Alzheimer's patient making a false claim about a nonexistent twin. It was indeed Tom who stole the book.

This seems to be a clear case of knowledge: You know that Tom stole a book at the library. And yet, according to the early version of the theory, you do not. There is, after all, a truth that is incoherent with your justification for believing that Tom is the thief, namely, D. Adding that truth to your belief system destroys your justification for the belief that P, since that truth clearly justifies you in believing a falsehood: that (F) the thief was actually Tom's indistinguishable twin. And, if F were added to your belief system, your justification for believing that P would obviously be destroyed. That's how a (non-actual) belief that D might be thought to defeat your justification for believing that P, even if F is not itself believed while D is believed. There is, counterfactually, a belief, namely, the belief that D, that justifies something the addition of which to your belief system destroys your justification for believing that P, namely, F. The mere commitment to F—your being entitled to believe it on the basis of D (if you believed that D)—is

[10] The defeasibilist's notion of truth-resistant justification is discussed at length in de Almeida [forthcoming]. The term 'truth-resistant justification' is originally from Moser [1989].

[11] At a crucial juncture in the development of the defeasibility theory, Swain [1974: 172–3] described the problem for Klein's and Hilpinen's early efforts as follows: 'Ironically, a man will sometimes wind up in [an epistemic] position that is worse than the one from which he began, even though he has moved closer to an ideally situated position by acquiring some new information.'

supposed to show that the belief that D destroys your justification for believing that P.[12] And, on the early defeasibility theory, if your justification for the belief that P does not withstand the addition of a truth to your belief system, the belief that P is Gettierized. But, patently, in DEMENTED, it is not, given that, as long as we allow for the inclusion of any other relevant truth, it becomes clear that your justification is not destroyed by the belief that Mrs. Grabit claims that Tom has a twin at the library. She's demented and there is no twin. This is the truth that R. Add both D and R to your belief system, and your justification for the belief that P remains unscathed. R is a *restorer* of your justification for believing—and *knowing*—that P.

Enter the so-called 'misleading defeaters': truths which do epistemic damage by justifying falsehoods which, in turn, if added to one's belief system, destroy any justification one may have for a given belief. And this is the reason for thinking that those truths are, as it were, malignant, and must have their malignancy neutralized. So, here's the twist: though true, a misleading defeater always works *through a falsehood*.[13] When a misleading defeater seems to destroy a justification, the damage is illusory: the impression of defeat depends on our failing to notice other, *restoring* truths in the vicinity, so to speak. In DEMENTED, the selective addition of the belief that D would give the impression that you're not justified in believing that P. But every misleading defeater is chased by a restorer, so to speak. The restorer is the 'defeater-eater' (in Klein's lingo). Which explains why the believer knows that P: don't stop at D while you're adding relevant truths to the would-be victim's belief system, or else you will have a distorted view of her epistemic situation.[14]

[12] More generally, the theory contends that an *initiating defeater* (e.g. our D) in a defeating (inferential) chain of propositions need not be the *effective defeater*. It is natural to think of an effective defeater as the contradictory of one of your beliefs, or as being what Pollock [1986] calls 'an undercutting defeater': a reason to think that that belief is somehow not justified for you. Initiating defeaters are like malignant tumours: they commit you to effective defeaters by justifying those defeaters for you (maybe in conjunction with some of your beliefs). This is at the core of our pre-theoretical understanding of what counterevidence is and of why a belief becomes untenable to us by making us *incoherent*. The defeasibility theory is, in a nutshell, the theory of epistemic defeat (loss of knowledge, or loss of epistemic justification) by incoherence (broadly construed, to include defeat by higher-order, *undercutting*, reasons).

[13] Klein's defeasibility theory is a precise account of this idea. See Klein [1981] and the technical point in note 14 below.

[14] More technically, the so-called 'misleading defeater' is the initiating defeater in a defeating chain that includes a falsehood. The initiating defeater must be true (as in DEMENTED). *But the effective defeater, the last link in the defeating chain, must be false*. Here's how we propose to establish this claim, thus plugging a gap in Klein's account of defeating chains. Suppose (for *reductio*) that (a) both the defeat is misleading (a pseudo-defeat) and the effective defeater (ED) is true. If the defeat is misleading, then, by hypothesis; (b) it works through a falsehood and the operative falsehood (F), the first falsehood in the chain, is not the initiating defeater (ID), since every ID is true; (c) there is a restorer (R, a defeater-eater) for the justification that implies the negation of the F; and (d) the conjunction F&R breaks the chain at or before its last link, the ED. If the ED is true, then, by (a), (c), and (d), the chain breaks before the ED. But, by (b), the chain must have an F, and the F is not the ID. Now, suppose that there is an F between the ID and the ED. In that case, given that the chain breaks before the ED, there is no chain to ED, simply because the chain is broken at the F. All talk of a chain that reaches to the ED through the F is misleading. (Recall this key principle of the theory: every justificatory chain, including defeating chains, must include only true links. That's the metaphor for saying that justification provided by a false link is not good enough for knowledge. The chain is broken by an R where the first falsehood turns up, but not before that first falsehood.) So, the F must be the ED. So, the ED is false. So, the ED is both true and false. Contradiction! So, either no defeat is misleading or the ED is the F. But there are pseudo-defeats, as illustrated by the DEMENTED case. Therefore, in every case of pseudo-defeat, the ED is false. Q.E.D. An obvious corollary of our (italicized) claim is that the justification relation between the ID and the ED cannot supervene on logical implication; it *must be* one of inductive support. We deem it important to see that Klein's defeasibility theory offers us the means to know exactly what defeating chains are made of. Although Klein has refrained from making the above (italicized) claim, he is entitled to make it. See Klein [1981: 142–8] for the discussion where the claim could have been made.

You have just survived a crash-landing in Gettierland! Now, brace yourself for the treacherous stretch of road ahead.

4. Guide to a Minefield in Gettierland

If you're collecting unanswered objections to the defeasibility theory, yours is a very small collection indeed! We believe that, by the end of December 2014, the discerning collector has gathered a total of three items in her collection, items contributed in turn by Feldman [2003], Foley [2012], and Turri [2012].[15] We shall now show you why these objections have reached their expiration dates.

Feldman thinks that the theory makes it hard for us to have a single belief whose justification remains undefeated. To see his point, consider how he expresses a *simple* defeasibility condition in an analysis of knowledge (one that is not constrained by the misleading/genuine defeat distinction). For Feldman [2003: 34],

> the [defeasibilist] proposal is to add to the [traditional analysis of knowledge] the requirement that there be no defeater:
>
> . . . S knows p = df. (i) S believes p; (ii) p is true; (iii) S is justified in believing p; (iv) There is no true proposition t such that, if S were justified in believing t, then S would not be justified in believing p. (No truth defeats S's justification for p.)

Before we consider Feldman's proposed refutation of this condition (iv) (a refutation that would render the mature theory—the one incorporating the notion of a misleading defeater—useless as well), let us note in passing that there is an insidious error in his formulation of this simple defeasibility condition. It may be tempting to think that, in its counterfactual role as an overrider of the justification in question—that is, as a truth that we add to S's belief system and which is such that, when thus added, overrides that justification—the defeater-turned-overrider, Feldman's t, must be a member of S's system of *justified* beliefs. Otherwise—one may think—the justification in question will not be overridden. It may seem that all overriders must be *justified* beliefs in order to exercise their overriding power, as it were. This is explicitly required by Feldman's condition (iv) above, and it's an important error.[16] If we made such a requirement, there couldn't be any *rebutting effective overriders*—that is, there couldn't be any overriders that negate a given (propositional) justifier.[17] To see that, suppose that justifier to be the false belief J. In that case, the proposition that ∼J is a genuine defeater of the justification based on J. However, if we require that ∼J be justified for S (to believe), then simply adding ∼J to a belief system containing J becomes

[15] These are *choice* items, of course, only the hardest ones we have found among those that await a reply.
[16] This point was originally made in de Almeida [forthcoming].
[17] Combining Klein's terminology with Pollock's would give us a 'rebutting/undercutting' distinction for 'overriders'. See note 12 above.

impossible—unless we admit that both J and ∼J can be justified *simultaneously* for the same agent, which most of us will not allow. But we are enjoined by defeasibilism to consider the result of *simply adding* the defeater to an existing justification and noticing how the defeater becomes an overrider of that justification. That's all well and good, because given, for instance, a 'single-link' justificatory chain that looks like this, 'J justifies P', we are expected to see how this becomes the case, 'the conjunction (J & ∼J) does not justify P.' But, for this picture to have its intended effect on us, we cannot suppose that the rebutting overrider, ∼J, is justified for S (to believe). No rebutting overrider can ever be justified *while doing its job as an overrider*.

Although the error does not play a crucial role in the context of Feldman's proposed refutation, we find it important enough to compromise the intelligibility of the theory.[18] With this point of clarification out of the way, we can now see how Feldman proposes to refute the defeasibility theory.

According to him, condition (iv) above is too strong, for excluding clear cases of knowledge. Here's the argument. Feldman [2003: 34] asks us to consider the following scenario:

> Smith is sitting in his study with his radio off and Smith knows that it is off. At the time, Classic Hits 101 is playing the great Neil Diamond's great song 'Girl, You'll Be a Woman Soon'. If Smith had the radio on and tuned to that station, Smith would hear the song and know that [the radio] is on.[19]

If we further add to this the assumption that Smith would not have learned about the song's being played unless he had the radio on, we have a problem, according to Feldman. The problem is that, in view of the fact that it is true that (Q) *Classic Hits 101 is now playing 'Girl, You'll Be a Woman Soon'*, Smith cannot satisfy the defeasibility condition (iv) and know that (P) *the radio is off*. Why would that be the case? Here's Feldman's explanation in his own words [ibid.]:

> In our example, if Smith were justified in believing [that Q], then he would have his radio on and he would hear the song. But if that were the case, then Smith would not be justified in believing that the radio is off. So condition (iv) is not satisfied. There is a true proposition [Q] such that if Smith were to be justified in believing it, then Smith would not be justified in believing [that P].

[18] Moreover, we expect you to agree that the mistake is interesting in itself. It's one that even prominent defeasibilists are inclined to make. For instance, after introducing the familiar defeasibility condition, Moser [1989: 255] writes: 'The analysis assumes that any true proposition, *T*, is added to S's evidence, *E*, in the sense that S is justified in believing that (*E* & *T*).' We see nothing in Moser's work indicating that he might have noticed that a rebutting overrider can never be justified for the believer in question. Nor do we see that Klein's work is immune to the error. Speaking about the issue in full generality, Klein [1986: 263] writes: 'Note that the overriding proposition could be, but need not be, justified for S.' *Undercutting* overriders, on the other hand, *can* be justified for S. These are 'higher-order' beliefs about how S fails to be justified in holding the target belief, or about how the target belief cannot be a case of knowledge for S in the circumstances.

[19] For the record, we also, like Feldman, hate the song!

Our task here is twofold. First, we want to understand whether Feldman's attack on condition (iv) is successful. And then we should ask ourselves if a successful refutation of a defeasibility view based on (iv) is a successful refutation of the *historical* defeasibility theory. To anticipate, we shall say 'yes' to the first question and 'no' to the second.

To see that Feldman's objection is successful against a view based on (iv), some reconstructive work is needed. As far as we can see, Feldman's *highly* compressed argument tacitly assumes that condition (iv) is motivated by the defeasibilist's tacit adherence to the following principle, which we may call a principle of 'counterfactual stability' (or 'CS') for epistemic justification:

> **CS**: S has an undefeated justification for believing that P only if there is no actual-world truth T (no '@-T') which is such that, in the nearest non-actual T-world where S believes that T, S is not justified in believing that P.[20] (Alternatively: S's justification for believing that P is undefeated only if every @-T fails to defeat S's justification in the nearest non-@ T-world where S believes that T.)

If CS faithfully captures the unstated principle in Feldman's objection, we should note that conjoining CS with the earlier stipulation about the only way in which Smith could have learned about the station's playing the song allows us to proceed with the reconstruction as follows. Schematically: In @, Smith's radio is off—and we want to say that he knows that it is off. Now, if condition (iv) is both satisfied *and implies CS*, there is no nearby world where the radio is on but Smith is justified in believing that it is not. But consider the nearest non-@ Q-world where Smith believes that Q, say, *w*, and notice that, as per Feldman's stipulation about how Smith would know about the song's being played, in *w* Smith's radio is on and tuned to Hits 101, and he's justified in believing that it is. So, let us assume that, in the nearest non-@ Q-world, Smith is not justified in believing that his radio is off. The nearest Q-world is one where he cannot know that the radio is off, since, there, his belief that Q justifies the belief that ∼P. Now, @ is like *w* in that @, too, is a Q-world. It's a Q-world where, by hypothesis, Smith does not believe that Q. But, in order to defeat a justification, a truth need not be believed by the agent whose justification is defeated by that truth. So, in no nearby Q-world can Smith have an undefeated justification for believing that the radio is off. And obviously there is no Q-world nearer @ than @ itself. So, in @ he cannot know that the radio is off. So, Feldman's condition (iv) excludes what intuitively counts as a *bona fide* case of knowledge.

How good is the objection? We are ready to agree that, in view of Feldman's objection, no defeasibility theory based on condition (iv) is tenable. Epistemic justification is just not counterfactually stable (as required by CS). But the question now is that of whether Klein should worry about the objection. And the answer to that is 'Not at all!' Here's why. According to

[20] Notice that Feldman would have us speak of a T-world where S *justifiably* believes that T. But, as we have seen, no such T-world can be the nearest non-@ world, since, in the nearest non-@ T-world where he believes that T, he also has beliefs that justify ∼T for him. Again, as we have seen, no rebutting overrider can be justified while doing the overriding. So, we read the objection charitably.

the *historical* defeasibility theory, epistemic defeat occurs if and only if a justification succumbs to the addition of an @-T to the agent's belief system *in T-worlds where that system is otherwise identical to her @-system*. But notice how references to non-@ T-worlds can be seriously misleading.[21] As we consider the radio case, we should ask ourselves how the @-truth that Q might be thought to destroy Smith's justification for believing that P. A moment's reflection will make it clear that Feldman's objection depends essentially on our focusing on the counterfactual situation in which Smith would be listening to the song (in *w*) and believing that Q, and that, in turn, would be providing him with a justification for believing that the radio was on. *That* justification—the justification he doesn't have in @, where, by hypothesis, he doesn't believe that Q—would then, in *w*, keep him from being justified in believing that P. By contrast, in clear cases of defeated justification such as LAMB, there is, in @, a truth, T, that, if added to the agent's belief system in some T-world where the agent's belief system is otherwise identical to her @-system, destroys the justification that the agent has in @. That's what we may call a failure of '@-truth-resistance'. Epistemic defeat, according to the historical theory, requires only that a justification fail to be @-truth-resistant. And we have now seen that *@-truth-resistance does not imply counterfactual stability* for epistemic justification. This is why Feldman's objection fails to refute the historical defeasibility theory. But the important lesson afforded by that objection is that epistemic justification is not counterfactually stable.

Richard Foley's recent objection may also be described as one according to which Klein's theory is too strong, but it presses the theory on another front. After noticing a 'striking similarity' [2012: 93] between his own account of knowledge and the defeasibility theory, Foley suggests that the latter has the following disadvantage vis-à-vis the former. Consider Klein's [1981] case of Loretta, the woman who may know that she owes the IRS $500 in spite of the presence of a defeater of questionable status in the situation. After carefully preparing her tax return, Loretta, out of an abundance of caution, asks her accountant to review her work. She believes that she owes the IRS $500, but she can use a confirmation from an expert.[22] In his haste, the accountant, meaning to reply that there were *no* errors made in Loretta's tax return, makes a damaging typo in his reply: 'your return contains errors', he writes. The question is this: before Loretta opens her mail and reads the misleading letter, does she know that she owes $500?

When the case was put forward, Klein [1981: 142–7] saw it as an ordinary case of knowledge where there is a misleading defeater—along the lines of the DEMENTED case. Foley agrees that that is a legitimate understanding of the case. But he believes that we can make trouble for the defeasibilist with a variation on the case. Suppose that the following counterfactual is

[21] We're claiming that a condition (iv) *that implies CS*—one that is based on that unwarranted principle—would make trouble for the defeasibility theory. The problem with Feldman's condition (iv) is *not at all* that it involves a subjunctive conditional. As we see it, pseudo-problems have been given currency by those who see subjunctive conditionals as troublesome in themselves for defeasibilists. See Hilpinen [1988] and Klein [1996] for relevant discussion of those pseudo-problems.

[22] Does Loretta *fully believe* that she owes $500, if she is still asking her accountant for assurances? That's an unfortunate aspect of the case, but we're not fussing over it.

true: If Loretta had opened her mail and been exposed to the misleading evidence, she would have called the accountant's office and failed to reach him, because he would have been away on vacation, and Loretta would then have talked to his supervisor, who would have assured her that the accountant is extremely careful, and the supervisor would have put her in touch with two of the accountant's co-workers, who would have given her even more evidence of his accuracy.

But beware: it's easy to miss Foley's point. You may miss the point if, for instance, you imagine that Loretta has already been exposed to the misleading evidence and lost her belief that she owes $500. Trivially, when she loses the belief, she then doesn't know. Foley's point is that, while Loretta still believes that she owes $500, the above counterfactual might still be true of her. What if it is true? This is the problem that Foley sees for the defeasibilist [2012: 98]:

> As the story is expanded in this direction, some listeners may be pulled away from the intuition that Loretta knows she owes $500. But if so, defeasibility theorists will need to explain why. What is the genuine defeater? Is there some additional factor in the revised story that turns the relevant truth here (the truth that the letter from the accountant contains the sentence 'Your return contains errors') from a misleading defeater into a genuine defeater, or is there perhaps another, more complicated truth that defeats her justification without justifying a falsehood?

This may look like a new objection, and the fact that it does so to Foley is ample justification for alarm. But, on close inspection, Foley's objection turns out to be an old one in a new guise. This is too much like Gilbert Harman's newspaper case to count as a new objection. We believe that, *mutatis mutandis*, Klein's response to Harman's case applies to Foley's refurbished Loretta case.

Recall that Harman's [1968] case was put forward as a Gettier case that supposedly revealed a 'social aspect' of knowledge. In Harman's original case, a civil rights leader is assassinated. The local, very reliable, newspaper runs the story of the assassination. Tom reads the paper and (supposedly) learns that the crime has been committed. But while he forms the true belief that the civil rights leader has been murdered, a cover-up is already under way, aimed at preventing widespread outrage. When Tom comes to believe that the crime has occurred, people around him had already fallen prey to the false denials in the media. He was, at that point, unique in his environment for ignoring the denials. There is every reason to think that, as soon as he hears about the denials, he, too, will lose the true belief about the assassination. While untouched by the massive cover-up, does he know that the crime has occurred?

Klein noticed that Harman's case was similar to the DEMENTED case. Here's how he reacted to it [1976: 810]:

> Harman believes that S does not know that the civil rights worker was assassinated … I am not certain about that, primarily because of the similarity

> between this case and the misleading-evidence Grabit case … Suppose that Mrs. Grabit had convinced others that she had twin sons and that for some reason various newspaper and radio reports tell of her imaginary progeny when reporting about the theft of the book … Does the fact that many other people believe what Mrs. Grabit said turn this into a case in which S fails to know? Would one false newspaper account … show that S does not know? Does the number of false newspaper stories make a difference?

The moral to be drawn from such considerations is that intuitions concerning some of the Gettier cases may vary, to the point that it becomes entirely legitimate to doubt whether those are *bona fide* Gettier cases. The fact that what some of us see as a *bona fide* Gettier case may legitimately be perceived as a pseudo-Gettier case by others is one of the most interesting data for epistemological theorizing. How does the defeasibility theory accommodate the datum? Klein has championed the view that the theory does just what we expect from a successful theory: It's flexible enough to accommodate the range of relevant intuitions. Here's how he supports the claim [ibid.: 811]:

> Those who share Harman's intuitions would believe that 'The newspaper and radio accounts deny that the civil rights leader was assassinated' defeats the justification by itself … Similarly, if many people shared Mrs. Grabit's delusion, that fact alone may be sufficient to defeat the justification [for believing that Tom is the thief]. That is, when a sufficient number of hitherto reliable sources … assert a given proposition which is in conflict with what has been justifiably believed, that by itself is enough to defeat the justification.

The contrasting, and equally legitimate, view is the one according to which those, like S, who, by sheer luck, haven't fallen prey to the misleading counterevidence do know after all, even if they are somehow unlike those who would otherwise be their epistemic peers. But this is a grey, sliding, area: You may legitimately hold that both (i) S knows that Tom Grabit stole a book and (ii) Tom (Harman's infrequent newspaper reader) does not know that the civil rights leader has been assassinated, depending on what, and how much of it, you add to the cases. It is a datum for epistemological theorizing that some Gettier cases inhabit the grey area. It is not optional for a successful theory to acknowledge the datum.

Now, back to the plight of careful Loretta. In Foley's scenario, intuitions may vary as to whether Loretta knows that she owes $500. The pull against a knowledge claim is indeed there, as noted by Foley. Some of us will see the truth that there is a letter from the (otherwise) reliable accountant mentioning (nonexistent) errors as a genuine defeater of the justification for believing that she owes $500. Others will be swayed by the apparent fact that the existence of the misleading letter justifies a claim of ignorance just because it is evidence for the falsehood that the accountant believes that there are errors in Loretta's tax return. One's genuine defeater may be another's misleading defeater in some carefully chosen cases. Paraphrasing the Randy Newman song, for the epistemologist, it's a jungle out there!

On close inspection, the defeasibilist has had a perfectly satisfactory reply to Foley's objection for nearly forty years now.[23]

5. The Perils of Turrism in Gettierland

Both of the objections we have seen claim that the defeasibility theory is too strong, for excluding either clear cases of knowledge (Feldman) or possible cases of ignorance (Foley). Turri's objection is even more ambitious. If Turri is right, the defeasibility theory has been a hopeless venture ever since the genuine/misleading-defeater distinction was introduced. But, given that there was no hope for the theory without that distinction, the theory is hopeless from inception.

According to Turri, the mature defeasibility theory—which he calls 'the modified defeasibility theory', or 'MDT', the one equipped with the notions of misleading defeaters and defeater-eaters (which, following Pollock [1986], he calls 'defeater defeaters')—suffers from a fatal flaw. Here's how Turri describes the problem [2012: 219]:

> Call this the *modified defeasibility theory*: (MDT) you know that P just in case (i) P is true, (ii) you believe in P based on evidence E, (iii) E justifies belief in P and (iv) E is ultimately undefeated. E is ultimately undefeated just in case there is no fact F such that (E+F) fails to justify belief in P; or if there is such a fact, then there is some further fact F^* such that (E+F+F^*) does justify belief in P . . . In such a case F^* is a defeater defeater . . . The main problem with MDT . . . is that the very device it introduces to give the intuitively correct verdict in [DEMENTED] also deprives it of the ability to handle the original Gettier cases. Consider LAMB. The fact that Linus is deceiving Dr. Lamb is a defeater (= F). But the fact that the modest female student does own a Lamborghini is a defeater defeater (= F^*). This last fact is a defeater defeater because this combination,
>
> E: My student Linus has possession of this Lamborghini, drives it frequently, and has a title to the Lamborghini with his name and birthdate on it; and
>
> F: My student Linus does not own this Lamborghini; and
>
> F^*: That young female student of mine owns a Lamborghini,
>
> justifies Dr. Lamb's belief that at least one of her students owns a Lamborghini. It does this because F^* obviously entails that at least one of her students owns a Lamborghini. And it would do so, no matter how many of Dr. Lamb's other students don't own a Lamborghini.[24]

Turri's MDT is a failed representation of the *historical* defeasibility theory, however. It omits a crucial element in the characterization of a defeater-

[23] The reader should contrast the defeasibility theory with its main rival, the sensitivity/safety account held by those who have followed Nozick's epistemology. See Nozick [1981] and Sosa [1999]. You will see that the Nozickian alternative is the one according to which neither Loretta nor the infrequent newspaper reader, Tom, is a knower in those cases, *and that's all there is to it*. No jungle out there!

[24] In what follows, 'F' and 'F^*' will be used to identify the relevant *facts*, as per Turri's stipulation.

eater, as we shall now see, and the omission accounts for his seemingly devastating objection.

One of Klein's invaluable contributions to the topic is an account of how the facts surrounding a true belief might mislead us into thinking that a justification that has *definitively* been destroyed remains restorable by the addition of true beliefs to the agent's doxastic system. As a result of the misimpression, one may mistakenly think of every defeated justification as a restorable one, as if there were a defeater-eater for every defeater. Turri's objection above is an acute form of the pathology.[25]

Consider the contrast between two of our cases above, one in which the historical defeasibility theory tells us that the believer is not Gettierized and one in which, *pace* Turri, the theory rightly gives us the result that the believer *is* Gettierized. In DEMENTED, there is a defeater-eater for the defeater of that justification. The reason for thinking that you know that Tom stole the book is that the defeater-eater—namely, the truth that Mrs. Grabit is demented and there is no twin—cancels the defeating effect of the defeater, the truth that Tom's mother speaks of Tom's being away and having an identical twin at the library. So, the justifying effect of the believer's *original evidence* is kept intact (or 'restored', as the metaphor would have it). In DEMENTED, the misleading defeater is clearly 'chased' by a defeater-eater. The believer *knows*, and we know how she knows.

But, in LAMB, the believer is obviously Gettierized—and that's exactly what the *historical* defeasibility theory tells us. As Klein teaches us, *sometimes a defeater creates a new justification while it destroys the old one*.[26] In LAMB, if we conjoin F with F*, as Turri does, the conjunction plays two roles: it destroys the old justification *while it creates a new one*.[27] There is no defeater-eater in the case! None. The old justification is hopelessly misleading. Linus is a con artist. Dr. Lamb has been had. She cannot know what she truly believes on that basis (Linus's misleading show of wealth). And, since she's Gettierized, the target belief is true. Had she enjoyed any access to the fact that makes it true, by having *evidence* to believe the relevant truth, namely, that the discreet student owns a Lamborghini, her belief that at least one of her students owns the exotic car would have been a case of knowledge. That is not what we have been asked to assume, however. We have been asked to assume that all of her evidence for the true belief comes from Linus's behaviour. And that evidence is *hopelessly misleading*. No truth can restore the justification based on it, even though some of the relevant truths in the circumstances would surely (virtually) *create* a *new* basis for knowledge-yielding justification.

[25] He would have found prophylactic treatment in Klein's discussion of John Barker's account of epistemic defeat. See Klein [1981: 151–7]. Klein [1980] rehearsed the point in a paper that is even more neglected than the corresponding passage in his [1981] book.

[26] See Klein [1981: 151–7].

[27] To be sure, there is some measure of inadequacy in Klein's talk of a defeater's 'creating a new justification'. The sympathetic reader will, however, understand that Klein cannot be speaking literally. If taken literally, the talk of a defeater's 'creating a new justification' is unintelligible since, by definition, a defeater is a true proposition *not in the believer's doxastic system*. So, literally speaking, no defeater can ever create a new justification for a given belief unless *it ceases to be a defeater and becomes a piece of evidence* for that belief by entering the believer's doxastic system. Charitable interpretation is therefore required for us to make sense of the notion that a defeater might *create a new justification* while destroying the old one. But it's definitely not a huge amount of charity!

That is the bare-bones version of our reply to Turri. But this is notoriously slippery ground. (Here we are, 35 long years after this crucial point was originally made still feeling the need to be reminded of it!) So, let us go over the point in slow motion, so to speak.

While reading the literature on the theory, especially the rushed, introductory, literature on the literature, one may, indeed, be tempted to reason as follows. 'According to the defeasibility theory, the epistemic agent knows that P only if the justification on which the agent's belief that P is based does not succumb to the addition of *any relevant truths* to the agent's belief system. In LAMB, the truth expressed by, say, "Linus is not a Lamborghini owner" seems to destroy the agent's justification for believing (the proposition expressed by the sentence) that (P) "There is a Lamborghini owner in the class." But if we add, to the agent's doxastic system, the obviously relevant truth that "The discreet young woman who sits in the back row owns a Lamborghini", we get the result that the agent is *still justified* in believing that P. So, according to the theory, the agent knows that P. Her justification for believing that P does survive the addition of relevant truths to the agent's doxastic system. However, it's obvious that Dr. Lamb *does not know* that P. Therefore, the defeasibility theory is false.'

That is exactly how Turri sees the issue. And it's a mistake. The mistake takes the form of what he calls the 'MDT'.[28] Turri's version of the defeasibility theory, his MDT, yields the above piece of reasoning. And the above piece of reasoning licenses his objection to what he takes to be the historical defeasibility theory. But, as we have now seen, that piece of reasoning is irrelevant to the historical theory, in view of the theory's plausible creation/restoration distinction. Turri's interpretation of the historical theory would imply that the agent's justification is not definitively destroyed. But it *is* definitively destroyed. And we are now in a position to see that the misinterpretation hinges on ignorance of the following theoretical point: A justification for which there is a defeater may not, however, be definitively destroyed. It will not be ultimately destroyed, only if it is *restored*. Restoration depends on the existence of *defeater-eaters*, truths that imply the negations of any would-be effective defeaters of the justification.[29] Any other truth on which the target belief may justifiably be based is a *justification-creator*, not a *justification-restorer* (not a defeater-eater). Thus, in LAMB, the truth that the young woman is a Lamborghini owner does absolutely nothing to restore Dr. Lamb's justification for the target belief, since it does not imply that Linus is not a Lamborghini owner. Dr. Lamb does not know that she has a Lamborghini owner among her students. The defeasibility theory explains why she fails to know. Turri's objection is unfair to the theory. It can only arise from inattention to the creation/restoration distinction.

[28] More specifically, the mistake resides in the simplicity of his characterization of ultimate defeat. As we have seen, some values of F* work by (virtually) *creating*, rather than *restoring*, a justification (in Klein's lingo). A clause making the creation/restoration distinction *must* be appended to the explanation of ultimate defeat. But this has been a tragic omission in the defeasibilist literature—Klein's more technical work being the exception.

[29] We should note, in passing, that we haven't been able to think of cases where a defeater-eater does not ultimately take the form of the negation of a would-be effective defeater. So, as far as we can see, every defeater-eater is a *rebutting* defeater-eater, as opposed to an *undercutting* one. We thank an anonymous referee for this journal for pressing us on the issue.

It is interesting to see how the distinction between the restoration of a justification and the creation of a new one by the very (conjunctive) defeater that defeats the old justification satisfactorily explains why the believer does not know in LUCKY LAMB, either. Given Dr. Lamb's evidence in the case, the defeater according to which Linus is a con artist conclusively defeats her justification for regarding him as a Lamborghini owner. Nothing, *absolutely nothing*, can restore *that* justification. Again, there are no defeater-eaters in sight. But, if you conjoin that defeater (that Linus is a deceiver, or 'F') with the truth that, despite all of his efforts to deceive, Linus is the rightful owner of the Lamborghini (or 'F*'), we have a new basis for knowing that he owns the car. But again, by hypothesis, no new evidence has been added to Dr. Lamb's belief system. A new justification is (virtually) *created* by the relevant conjunction (F&F*). But Lamb is not lucky enough to have it.

6. Concluding Remarks: A Timely Reminder

It is disconcerting to us to see how very little attention has lately been paid to the theory that dominated epistemological theorizing in the 1970s, the defeasibility theory of knowledge. We have not attempted to write the story of its fall from grace. This is not the place for a long report on the theory's present state of social neglect. But we'll give you this one startling piece of evidence for the claim. If you look at the *Stanford Encyclopedia of Philosophy* entry on 'The Analysis of Knowledge', by Jonathan Jenkins Ichikawa and Matthias Steup, in December 2014, you will find a chunky discussion of reactions to the Gettier Problem, but you won't even see the *word* 'defeasibility', much less an introduction to the theory, or even just a few words reporting on its historical significance. The name 'Klein' cannot be found in the entry, either.

Again, we're not pursuing the non-philosophical aspects of trendsetting in epistemology, or the causes of community-wide memory loss. We have looked into the possibility that actual *philosophical failure* may have been a factor in the theory's dramatic loss in popularity. And we haven't found good reason for dismissing the theory in the objections recently put forward by three influential philosophers.

We have also suggested that Turri's objection is the most dangerous of the bunch. As we see it, defeasibility theorists have left the door open to an important, exciting, error such as Turri's, at the very least by failing to emphasize a crucial aspect of the theory. And we submit that one who does not clearly see why Turri's objection is unfair to the defeasibility theory is precariously situated to hold any views on whether the Gettier Problem has been solved. Most will say that it has not been solved. Those voices are legion. And they may be right. The claim that the Gettier Problem has not been solved is, as a matter of sociological fact, one of the safest claims one can make in the philosophical community in an effort to protect one's reputation. Hordes of undergraduate students are instructed by it. This paper is not a refutation of that claim. But it is a plea for more caution in 'Gettierland', much more caution than the hordes are getting with their

education in epistemology. The defeasibility theory may not ultimately prove to be the last word on the Gettier Problem. But we submit that it remains a formidable contender in that fight. For some of us, there is no persuasive objection to it in sight, apart perhaps from the one that motivates the recent developments in Klein's [2008] work on the theory.[30] And we think that understanding why the three objections tackled here cannot reasonably be used to dismiss the theory is not merely optional for anybody aiming to understand the current state of the debate over the Gettier Problem.[31]

PUCRS (Brazil)

Received: May 2014
Revised: January 2015

References

Chisholm, Roderick M. 1964. The Ethics of Requirement, *American Philosophical Quarterly* 1/2: 147–53.
de Almeida, Claudio [forthcoming]. Knowledge, Benign Falsehoods, and the Gettier Problem, in *Explaining Knowledge: New Essays on the Gettier Problem*, ed. Rodrigo Borges, Claudio de Almeida and Peter D. Klein, Oxford: Oxford University Press.
Feldman, Richard 1974. An Alleged Defect in Gettier Counter-Examples, *Australasian Journal of Philosophy* 52/1: 68–9.
Feldman, Richard 2003. *Epistemology*, Upper Saddle River, NJ: Prentice Hall.
Foley, Richard 2012. *When Is True Belief Knowledge?* Princeton: Princeton University Press.
Harman, Gilbert 1968. Knowledge, Inference, and Explanation. *American Philosophical Quarterly* 5/3: 164–73.
Hetherington, Stephen 2012. The Gettier-Illusion: Gettier-Partialism and Infallibilism, *Synthese* 188: 217–30.
Hilpinen, Risto 1988. Knowledge and Conditionals, in *Philosophical Perspectives, 2, Epistemology*, ed. James Tomberlin, Atascadero, CA: Ridgeview: 157–82.
Hintikka, Jaakko 1962. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, new edition, ed. Vincent F. Hendricks and John Symons, London: King's College Publications, 2005.
Ichikawa, Jonathan Jenkins and Steup, Matthias 2014. The Analysis of Knowledge, in *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), ed. Edward N. Zalta, URL = <http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>.
Klein, Peter 1971. A Proposed Definition of Propositional Knowledge, *The Journal of Philosophy* 68: 471–82.
Klein, Peter D. 1976. Knowledge, Causality, and Defeasibility, *The Journal of Philosophy* 73: 792–812.
Klein, Peter D. 1980. Misleading Evidence and the Restoration of Justification, *Philosophical Studies* 37: 81–9.
Klein, Peter D. 1981. *Certainty: A Refutation of Scepticism*, Minneapolis: University of Minnesota Press.
Klein, Peter D. 1986. Immune Belief Systems, *Philosophical Topics* 14/1: 259–80.
Klein, Peter D. 1996. Warrant, Proper Function, Reliabilism, and Defeasibility, in *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, ed. Jonathan L. Kvanvig, Lanham, MD: Rowman & Littlefield: 97–130.
Klein, Peter D. 2008. Useful False Beliefs, in *Epistemology: New Essays*, ed. Quentin Smith, Oxford: Oxford University Press: 25–61.
Kvanvig, Jonathan L. 2005. A Perspective on Plantinga's Theory of Warrant, *Certain Doubts*, post # 308, http://certaindoubts.com/?p=308.
Lehrer, Keith 1965. Knowledge, Truth and Evidence, *Analysis* 25/5: 168–75.
Lehrer, Keith and Paxson, Thomas 1969. Knowledge: Undefeated Justified True Belief, *The Journal of Philosophy* 66/8: 225–37.
Lycan, William G. 2006. On the Gettier Problem Problem, in *Epistemology Futures*, ed. Stephen Hetherington, Oxford: Clarendon Press: 148–68.
Moser, Paul K. 1989. *Knowledge and Evidence*, Cambridge: Cambridge University Press.
Nozick, Robert 1981. *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
Plantinga, Alvin 1996. Respondeo, in *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, ed. Jonathan L. Kvanvig, Lanham, MD: Rowman & Littlefield: 307–78.

[30] This is not to ignore the recent option, most notably represented in Hetherington [2012], according to which the Gettier Problem as a whole, including the proposed solutions that play by its principles, will eventually be seen as a pseudo-problem. For discussion of other influential instances of opposition to the Gettier Problem as a whole, see Lycan [2006].

[31] We thank Rodrigo Borges, Peter Klein, and two anonymous referees for the AJP for helpful discussion of versions of this paper. We should also like to express our gratitude to the Brazilian funding agency CNPq for the PIBIC scholarship that partially funded J.R. Fett's work on this project.

Pollock, John L. 1986. *Contemporary Theories of Knowledge*, Totowa, NJ: Rowman & Littlefield.
Shope, Robert 2002. Conditions and Analyses of Knowing, in *The Oxford Handbook of Epistemology*, ed. Paul K. Moser, Oxford: Oxford University Press: 25–70.
Shope, Robert 2004. The Analysis of Knowing. In *Handbook of Epistemology*, ed. Ilkka Niniluoto, Matti Sintonen, and Jan Wolenski, Dordrecht: Kluwer: 283–329.
Sosa, Ernest 1999. How to Defeat Opposition to Moore, in *Philosophical Perspectives, 13, Epistemology*, ed. James Tomberlin, Oxford: Blackwell Publishers: 141–53.
Swain, Marshall 1974. Epistemic Defeasibility, *American Philosophical Quarterly* 11/1: 15–25.
Swain, Marshall 1981. *Reasons and Knowledge*, Ithaca, NY: Cornell University Press.
Swain, Marshall 1996. Warrant versus Indefeasible Justification, in *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, ed. Jonathan L. Kvanvig, Lanham, MD: Rowman & Littlefield: 131–46.
Turri, John 2012. In Gettier's Wake, in *Epistemology: The Key Thinkers*, ed. Stephen Hetherington, London: Continuum: 214–29.