

“There Is No Knowledge from Falsehood”
DRAFT: Please cite published version (forthcoming in *Episteme*)

By Ian Schnee
ian.schnee@wku.edu

1. Introduction

A growing number of authors defend putative examples of knowledge from falsehood (KFF), inferential knowledge based in a critical or essential way on false premises, and they argue that KFF has important implications for many areas of epistemology (whether evidence can be false, the Gettier debate, defeasibility theories of knowledge, etc.).¹ I argue, however, that there is no KFF (or, at least, there is no compulsion to accept KFF), because in any supposed example either the falsehood does not contribute to the knowledge or the subject lacks knowledge. All the recent arguments that rely on KFF are therefore based on a mistake.

Here is the outline. In Section 2 I clarify the nature of KFF and the range of cases in which it supposedly appears. In Section 3 I argue that, contrary to the emerging view, it is more plausible to maintain that subjects in the cases lack knowledge than to maintain that they possess knowledge. In Section 4 I consider KFF cases presented by Fitelson (2010), in which the knowledge is counterfactually dependent on the falsehood (which Fitelson calls KFF*), and I show that the argument in Section 3 works for KFF* cases too. Thus, regardless of the case, be it traditional KFF or KFF*, there is no knowledge from falsehood. Finally, in Section 5 I present a positive account (the Ersatz Solution) of how knowledge works in the relevant cases and respond to objections that KFF advocates have presented.

¹ Warfield 2005; Klein 2008; Fitelson 2010; Arnold 2013; cf. Hilpinen 1988, Klein 1996, Luzzi 2010, and Rizzieri 2011.

2. Knowledge from Falsehood Clarified

Subjects can make inferences from specific but false information. Consider:

TV Show 1: Ellie's favorite TV show is on from six to seven. Ellie looks at her watch, which reads six thirty, and she infers, from the consideration that it is (exactly) six thirty, that her show is on. But Ellie is in error. It isn't (exactly) six thirty, it is six thirty-two; her watch is slow by two minutes.²

Despite her error it seems that Ellie knows that her favorite show is on. That might be surprising, since she made an inference from a false premise and thus it seems that her knowledge depends in an important or essential way on falsehood.

I start with two clarifications about the example. First, I am generally neutral about the metaphysical status of Ellie's evidence or reasons. I will speak of propositions (and facts) being her evidence, but my points can be recast in terms of the mental states whose contents are those propositions. I maintain that we should reject KFF whether evidence and reasons are propositions, facts, mental states, or anything else. Second, I distinguish as a conceptual point between a piece of evidence that a subject uses in a conscious inference and the evidential basis of her knowledge (the grounds on which she has inferential knowledge); call the first notion S-

² To clarify our intuitions we can also stipulate that her watch is very reliable about the approximate time and Ellie knows as much. For similar and additional examples see Warfield (2005), Klein (2008), Arnold (2013), etc. The qualification "(exactly)" is included in order put aside one worry about this and several other KFF cases: one might think that without it Ellie's belief that it is six thirty would be true, because there is a margin of error (perhaps context dependent) that affects the truth of what she believes (cf. Travis 1997). I will assume that these sorts of worries can be dealt with by KFF advocates.

evidence and the second notion **K-evidence**.³ In a typical case one proposition is both a subject's S-evidence and K-evidence, but the two notions are not conceptually or analytically guaranteed to coincide. It is clearly true that Ellie's S-evidence in TV Show 1 is false (she makes a conscious inference from the proposition that it is six thirty, which is false), and thus one can have knowledge through an inference that uses false S-evidence.

The KFF advocates, however, universally claim something stronger: they claim that K-evidence can be false.⁴ By knowledge from falsehood, I will thus mean inferential knowledge from false K-evidence. TV Show 1, though, does not show that K-evidence can be false without supplementation, e.g., by

Sufficiency Thesis: If subject *S* has inferential knowledge that *p* via inference from S-evidence *e*, then *e* is also K-evidence for *S*.

Sufficiency says that all S-evidence is K-evidence; given the fact that Ellie has false S-evidence, it would follow that she also has false K-evidence. Thus there would be KFF.

³ Cf. Warfield (2005: 409–410) and Arnold (2013: 170). KFF concerns *inferential* knowledge from falsehood, and thus I assume there must be some grounds or evidence on the basis of which the subject has knowledge (additional discussion in footnote 8).

⁴ Examples: Warfield (2005: 408–413); Klein, given some interpretive issues (2008: 25–26); Fitelson (2010: 668–669); Arnold (2013: 171). Notice that if KFF only concerned false S-evidence then it couldn't do any of the interesting epistemological work one would hope it could, since the relevant evidential notions employed by theories that KFF is supposed to bear on (e.g., Williamson's (2000) E=K thesis) is K-evidence (see Warfield 2005: 405, Arnold 2013: 165). The important exception seems to be Littlejohn (2013), who is willing to accept, at least on some understanding, KFF and thus might be considered a KFF advocate; Littlejohn's (2013) position is different from traditional KFF advocates, however, and I discuss it in footnote 8 below.

No KFF advocate has argued for Sufficiency, however.⁵ Indeed, one might naturally think that it has counterexamples. For example, consider this case, due originally to Lehrer:

Suppose that I am doxastically justified in believing that Havit owns a Ford (which is true) and also justified in believing that Nogot owns a Ford (which is false). On the basis of those two beliefs, I infer, and thereby come to know, that someone in the class owns a Ford.⁶

The subject consciously reflected on two premises when he made his inference, so there are two pieces of S-evidence, one true and one false. We might think, though, that the false S-evidence (about Nogot) is irrelevant to his knowledge—his knowledge actually derives from the true S-evidence (about Havit). I will not argue here that that is how we must understand the case; I am just illustrating why one might think that there are counterexamples to Sufficiency: the Nogot S-evidence is seemingly not K-evidence.⁷

Instead of using Sufficiency, KFF advocates' strategy has been to alter examples like TV Show 1 in a key way:

⁵ Arnold might appear to be an exception at 2013: 167, but he is not; see 2013: 170–171.

⁶ I am quoting Klein's (2008: 41) rendition of the case. See Lehrer (1990: 137–138) for a slightly different presentation of the same basic idea.

⁷ KFF advocates appear to agree. Klein, e.g., calls the false S-evidence in Lehrer's example a "harmless falsehood", which he contrasts with the "useful falsehoods" of KFF proper (Klein 2008: 42–42; also see Warfield 2005: 405–406). Concurrence, noted by Klein, can also be found in Lehrer (1990: 138) and Goldman (1967: 368).

TV Show 2: This case is like TV Show 1, but Ellie is extremely confident about the accuracy of her watch and only forms exact beliefs from it; she has no other evidence at all regarding the time.

The crucial difference is that TV Show 2 eliminates all other potential evidential grounds for Ellie's knowledge. Thus if she has inferential knowledge, her S-evidence must be her K-evidence; there is simply no other candidate. KFF advocates can therefore avoid Sufficiency and instead use a weaker principle:

Solitary Thesis: If subject *S* has inferential knowledge that *p* via inference from S-evidence *e*, and *e* is the sole evidential path that could (given various facts about the case) enable *S*'s knowledge, then *e* is also K-evidence for *S*.

Solitary is much harder to reject than Sufficiency (recall that our concern is inferential knowledge, so I am granting that there must be a knowledge-enabling evidential path).⁸ Thus if

⁸ An important challenge to Solitary is Littlejohn (2013), who accepts that there can be inferential knowledge based on/from a false premise but denies that the false premise constitutes evidence. Littlejohn's reasons for rejecting Solitary, however, are that it conflicts with endorsing two other claims, namely (i) KFF (in the sense that inferential knowledge can be gained by inference from a false premise) and (ii) evidence truthism (all evidence is true). He does not give independent reasons for rejecting Solitary. Littlejohn carefully motivates evidence truthism, but it is important to note that this rejection of Solitary is consequently dependent on the prospects of that thesis. More importantly, however, if my arguments against KFF are correct then this rejection of Solitary is still unmotivated; evidence truthers simply have no reason to incur the cost of rejecting such a plausible principle as Solitary (the Ersatz Solution that I defend in Section 5 is compatible with evidence truthism). To the point: the rejection of Solitary faces the following dilemma. One who rejects Solitary must hold that knowledge in putative KFF cases is either genuinely epistemically inferential or not (such as merely causally inferential, in the sense that there was a causal process of reasoning from premises that led to the belief, regardless of whether that process constitutes an evidential pathway responsible for the belief's epistemic

Ellie has knowledge in TV Show 2, there is genuine KFF. Solitary also provides a helpful way of understanding how a falsehood could be *essential* to knowledge—if the falsehood is K-evidence that figures in the sole evidential path enabling *S*'s knowledge, then there is a clear sense in which the falsehood is essential. I will therefore understand KFF in general to be any case in which a subject has knowledge grounded on false K-evidence, and essential KFF to be a case of KFF in which the falsehood is part of the sole evidential path enabling the subject's knowledge. A non-essential case of KFF would then be a case in which a falsehood is a genuine

credentials). Say one takes the first case; the belief is genuinely epistemically inferential. Then there must be some K-evidence that grounds the knowledge. After all, K-evidence, as I understand it, just is the grounds of one's epistemically inferential knowledge (I do not need to presume that evidentialism is true). But then Solitary follows necessarily (holding that the K-evidence need not enable the knowledge would lead to obvious basing objections). That leaves the second case—one must hold, then, that KFF cases are actually epistemically non-inferential knowledge. As a side note, this is not how KFF advocates have understood their cases; even Klein (2008) seems to admit that inferential knowledge from falsehood cases must have K-evidence (as does Warfield, Arnold, etc.). But there are two more important points. First, it is dubious to hold that these cases are epistemically non-inferential; Ellie's knowledge in TV Show 1 that her show is on does not seem to be epistemically basic or non-inferential (we could also construct an analogous good case in which the knowledge would certainly be inferential, even by evidence-truthers' lights). Second, it is even more problematic to hold that these cases are non-inferential and yet *from* the falsehood in any epistemologically significant way. To highlight the point, consider a case where one has two independent epistemic routes to a belief, one basic and one inferential (assume both are sustaining causes or explanatory for basing, but the inferential path was the proximate cause (S-evidence)). If the inferential path is via a falsehood, one might maintain that the belief is knowledge because of the basic path. So even though the proximate cause was an inference, the belief *qua* knowledge is genuinely epistemically basic. But this is not the model the Solitary rejecter is asking us to accept, since this is not KFF at all. This knowledge is properly basic and is completely independent of the falsehood (compare the Havit and Nogot case above regarding Sufficiency; for elaboration see Section 5). The Solitary rejecter is trying to have it both ways (trying to make the knowledge from the falsehood, and thus from a process of reasoning from a premise, and yet epistemically non-inferential), with implausible consequences. I therefore find the rejection of Solitary unmotivated and problematic—evidence truthers should endorse the Ersatz Solution, not reject Solitary. Thanks to an anonymous reviewer for pressing me on these issues.

bit of K-evidence for the subject, but in which the evidential path involving that falsehood is not the sole evidential path enabling the subject's knowledge.⁹

In the next section I argue for two claims: (1) essential KFF, so understood, does not exist; (2) if there is no essential KFF then there is no KFF at all. Before doing so, however, I want to focus on what is required in order for S-evidence *e* to be the sole evidential path that could enable *S*'s knowledge. An evidential path, in the simplest case, is a sequence of propositions, each of which bears evidential support to the next.¹⁰ In order to enable knowledge for a subject, however, an evidential path must meet further conditions. For example, I will grant that for each proposition in the path:

- (i) the subject must believe the proposition, and
- (ii) the beliefs from (i) must bear the right causal or explanatory relations (or whatever the basing relation requires) to the subject's knowledge.¹¹

An essential KFF case, then, is a case in which there is no independent evidential path that meets these two conditions. For illustration, let me clarify why one might think that TV Show 1 is not essential KFF. Even though in TV Show 1 Ellie's belief that it is (exactly) six thirty is false

⁹ If, that is, there could be two independent evidential paths which actually enable a subject's knowledge. Whether that is so I do not try to adjudicate; I merely leave the conceptual room for it.

¹⁰ For more on evidential paths see Klein (2003, 2008).

¹¹ These beliefs can be dispositional beliefs, because dispositional beliefs can stand, I will argue in Section 3, in the right causal or explanatory relations. By dispositional belief I mean a genuine but non-occurrent belief, which I distinguish from a mere disposition to form a belief (for discussion see Schwitzgebel (2011); thanks to Peter Murphy here). This distinction is no liability for me: without this distinction it is even harder to make the case for KFF. I do not claim (i) and (ii) are sufficient. There are additional necessary conditions as well. E.g., (iii) the level of evidential support cannot just be some sort of weak confirmation, but must be sufficiently strong to enable knowledge.

(because it is six thirty-two), it is compatible with the example that Ellie also believes that it is *approximately* six thirty (cf. Warfield 2005: 410) and that, barring further argument to the contrary, her belief in the approximate time bears the right causal or explanatory relations to her knowledge. There is therefore an independent evidential path in TV Show 1, a path which contains no falsehood. Thus even if Ellie's false belief is K-evidence for her, it would not be essentially false K-evidence. Furthermore, if it turns out that the evidential path involving the approximate belief is the sole source of her knowledge, to the exclusion of the falsehood-involving path, then the case would not be a case of KFF at all (her S-evidence would not be K-evidence).

3. Two Arguments

3.1 Essential KFF

In this section I argue first against essential KFF and then against KFF in general. Essential KFF, recall, occurs when the falsehood is part of the sole evidential path enabling the subject's knowledge. Advocates of KFF all argue for KFF by presenting putative cases of essential KFF and claiming that intuitively the subjects in the cases have knowledge. Here was our example:

TV Show 2: This case is like TV Show 1, but Ellie is extremely confident about the accuracy of her watch and only forms exact beliefs from it; she has no other evidence at all regarding the time.

TV Show 2 blocks alternative evidential paths for her knowledge—e.g., it blocks the path involving the belief that it is approximately six thirty. So if indeed there is intuitive pull to the thought that Ellie in TV Show 2 has knowledge, then there is some intuitive pull to essential KFF. I will first make a methodological point about “intuitions”, since that term has received increasing philosophical discussion.¹² I am not interested in intuitions in the sense of mere gut reactions to cases; instead I am interested in considered judgments about cases based on an understanding of relevant epistemic features. Intuitions, as I will use them, can therefore be revised in light of better understanding, and revised intuitions, when they are based on a better epistemological understanding, trump unrevised ones.¹³ To the point: perhaps we have an initial intuition that Ellie in TV Show 2 has knowledge, but I will argue that once we understand the epistemic features of the case there is no intuitive pull to the thought that she has knowledge; on the contrary, there are good grounds for denying that she has knowledge.

The argument against essential KFF has two parts, because TV Show 2 is underspecified. Ellie is an unusual subject, because normally when we form exact beliefs about the time from a watch (or any other device), we also form, at least dispositionally¹⁴, the belief that it is approximately that time. The issue to focus on is whether or not Ellie has a good epistemic reason for her unusual doxastic behavior.

3.1.1 *Essential KFF, Part 1*

¹² A small sampling: Bengson (2013), Cappelen (2012), and Williamson (2007).

¹³ I take this to be at least one fairly common usage in epistemology, but I do not claim that this is the only way intuitions could or should be used in epistemology or philosophy in general (cf. Pust 2012).

¹⁴ As noted in footnote 11, I distinguish between a dispositional belief and a mere disposition to believe. I take it that we often implicitly form dispositional beliefs without making conscious judgments.

Let's say she does have a good epistemic reason for not forming approximate beliefs.

For example, say there are facts that she knows which account for her behavior:

TV Show 3: This case is like TV Show 2, but Ellie knows that, even though her watch is quite reliable, when her watch is not exactly right it is usually an hour or more off. That is why she does not form approximate beliefs about the time from her watch; she believes exactly what her watch says or she forms no belief at all about the time on its basis.¹⁵

Once we fully understand TV Show 3, it is clear that Ellie does not know that her show is on. She is in a Gettier case (Gettier 1963); her belief is problematically lucky or accidental. First, she is the victim of bad luck, because her normally reliable watch is malfunctioning in this case; but she is also the victim of good luck, because when her watch does malfunction it is usually off by an hour or more, whereas right now it is only off by two minutes. She is thus problematically lucky that her inference to the belief that her show is on happens to lead her to the truth.¹⁶ The point generalizes: if there are facts that a subject knows that underwrite her avoidance of all alternative evidential paths—e.g., avoidance of the belief that it is approximately six thirty—then her inference will always be problematically lucky.

¹⁵ Thanks here to Martin Montminy, who provided helpful comments on earlier drafts of this paper. Montminy (ms.) discusses a similar type of example, and his work and feedback helped me greatly.

¹⁶ There is probably no uncontroversial characterization of Gettier cases, but I take it that a basic idea is that subjects in Gettier cases are problematically lucky in a certain way, and dual-luck presentations are one way of illustrating the problematic luck. For a dual-luck discussion of Gettier cases see Zagzebski (1994). On various types of epistemic luck see Pritchard (2005). (A different sort of issue is whether theories should be constrained by traditional judgments of Gettier cases because of potential cultural relativity, etc. (see Weinberg et al. 2001), but that issue is not the focus of the KFF debate.)

Before continuing to part 2 let me consider an objection. In TV Show 3 there are facts about the watch that underwrite the dual-luck Gettier analysis. The worry is this: what if there are no such facts, but she has a justified belief that there are? Then there won't actually be facts that make her belief succumb to dual luck in the way I described, but she will still have good epistemic reason for only believing exactly what her watch says. For example:

TV Show 4: This case is like TV Show 3, except that instead of knowledge, Ellie has a justified false belief that her watch is normally reliable but when it malfunctions it is off by an hour or more. She had a watch-testing machine run a test on her watch to determine exactly how it operates. The normally reliable machine, however, mixed up the reports and gave her the report for a different watch; hence her justified false belief. Her watch is actually like most watches; it is very reliable about the approximate time.

In TV Show 4 Ellie's inference is not problematically lucky in the way that I depicted in TV Show 3, because now being off by two minutes does not necessarily constitute a malfunction. I have therefore not yet shown that Ellie lacks knowledge in TV Show 4. There are still grounds, however, for thinking that she is in a Gettier case. Ellie's method of forming beliefs from her watch in this version essentially involves the machine's report, and so does her problematic luck: she is the victim of bad luck, because the normally reliable machine is wrong about how her watch operates and under what conditions the watch is reliable, but she is the victim of good luck because the report that she did get, despite the machine's error, happens to still lead her to the

truth. Once we understand the relevant epistemic features of the case, then, there are good grounds for thinking that Ellie lacks knowledge.¹⁷

Some readers may still be undecided about the cases at issue (TV Show 3 and 4). It is sometimes hard to make judgments, even on reflection, about odd cases. I have an additional reply for readers in this situation: the burden of proof is on advocates of KFF for two reasons. First, part of the growing interest in KFF derives from its bearing on significant issues in epistemology, such as Williamson's E=K thesis or the weaker thesis that all evidence is true or all reasons are facts. If the case for KFF is a wash, however, then it will be dialectically useless. Second, there is an initial presumption against KFF, because knowledge cannot be accidental or lucky (in a way that is difficult or impossible analyze) and if knowledge is essentially based on falsehood then there is a sense in which it seems problematically accidental or lucky. For example, consider a standard Gettier case: Max looks in the yard and sees what he takes to be a dog (though it is just a plastic statue); he believes that there is a dog in the yard and infers that there is an animal in the yard; unbeknownst to him, there is a squirrel in the yard (Warfield 2005). On the standard way of thinking, Max has a justified true belief that there is an animal in the yard, but he doesn't know it. History has taught us that it is hard to specify exactly what is

¹⁷ My contention is that Ellie's method of following her watch is Gettiered, and thereby her belief is Gettiered; cf. Sainsbury (1996) and Weatherson (2004). The machine instantiates a classic Gettier pattern: it is normally reliable, but happens to be wrong this time. We might try to eliminate that aspect of the example. One way to do so is to imagine that the machine isn't actually reliable; Ellie merely has a justified belief that it is. This move is an iteration of the move we already made, however, from TV Show 3 to 4. If that move didn't work the first time, it won't work the second. Alternatively, one might try to weaken Ellie's justification for her belief about the watch. E.g., say that instead of using the machine, Ellie just has fairly good inductive grounds for her belief about the watch's behavior. The problem, however, is that such a fix still does not eliminate the problematic luck. Ellie is the victim of bad luck, because her inductive grounds are misleading (they falsely suggest that her watch behaves in a certain way, which she relies on), but she is the victim of good luck, because, even though her grounds are misleading, she is nonetheless led to the truth.

wrong with Max's situation, but on one way of thinking about it Max doesn't understand why his belief is true, because his reasons don't connect his belief to the truth in the right way. The reason why his belief is true is that the squirrel is in the yard, not that there is a dog in the yard (as he supposes). Similarly, the reason why Ellie's belief is true is that it is approximately six thirty, not that it is exactly six thirty (as she supposes). I do not claim to have an analysis of Gettier cases¹⁸, so I cannot show that these suggestive similarities between Max and Ellie are conclusive; rather, my point is that if KFF advocates want to convince us that Ellie indeed has knowledge, they need clear and compelling cases. If TV Show 4 is simply unclear, then that is still a victory against KFF.

My assessment of part 1, then, is this. If there is a rational epistemic explanation for why Ellie does not hold the approximate-time belief (or any other belief that would enable an alternate evidential path), then insofar as we have clear intuitions about the cases those intuitions favor the verdict that she lacks knowledge. So far it is therefore more plausible to reject than accept essential KFF.

3.1.2 *Essential KFF, Part 2*

Next I turn to part 2: suppose that there is no rational epistemic explanation for why Ellie does not hold the approximate-time belief (or any other belief that would enable an alternate evidential path). For example:

¹⁸ In particular, I do not claim that the remarks above provide necessary and sufficient conditions for being in a Gettier case. Max and Ellie both essentially reason from a falsehood, so my remarks are only meant to characterize a *prima facie* similarity between them. Fake-barn county cases, e.g., do not seem to conform to my remarks, but those cases are also not Gettier cases that essentially rely on reasoning from a falsehood.

TV Show 5: This case is like TV Show 2; Ellie is extremely confident about the accuracy of her watch and only forms exact beliefs from it, but she has no reason for this doxastic behavior; she has no other evidence at all regarding the time.

There is an initial problem, however, with TV Show 5. There is something problematic about Ellie's method of belief formation, which gives us grounds for denying that she has knowledge. There is a mismatch between her confidence and her watch's accuracy: she is extremely confident about the accuracy of her watch, but her watch is not extremely accurate. Here we face another dilemma: her confidence is either based on a good epistemic reason or it isn't. If it is based on a good reason, then it is problematically lucky. She suffers from bad luck, because her good reasons to be confident in the extreme accuracy of her watch are misleading (her watch is not extremely accurate); but she suffers from good luck, because that mismatch does not lead her to error in this case; her belief is true despite the mismatch. If it is not based on any good reason, however, it shouldn't be surprising that it can fail to lead to knowledge: Ellie is extremely confident for no good reason in the accuracy of a watch that is not extremely accurate and refuses to form anything but exact beliefs on its basis. That situation is clearly epistemically defective. As such it gives us grounds to resist attributing knowledge to her.

This is not a knockdown argument that Ellie must lack knowledge in TV Show 5. In moving from TV Show 4 to 5, KFF advocates have given us less information about Ellie to work with. Although such a move makes it harder to demonstrate exactly why she lacks knowledge, it also dampens the force of any judgment that she has knowledge. So rather than a knockdown argument that Ellie lacks knowledge, I take the remarks above to show that, no matter how KFF advocates construct their cases, there is always something epistemically problematic about the

subject that gives us as good or better grounds for resisting the attribution of knowledge than for granting it. Given the burden of proof considerations raised in Section 3.1.1, that means that the case against KFF is stronger than the case for it.¹⁹

The lesson one might draw from TV Show 5 is that the way that KFF advocates have constructed their cases is problematic. In order to block all alternate evidential paths, KFF advocates have supposed that subjects like Ellie have extreme confidence that, on reflection, creates epistemic problems.²⁰ Perhaps the solution, then, is to drop that aspect of the example.

Consider:

TV Show 6: This case is like TV Show 1, but we just stipulate that Ellie only forms exact beliefs from her watch, so that she lacks beliefs from her watch in the approximate time or in any other propositions that could constitute an independent evidential path; she has no other evidence at all regarding the time.

¹⁹ It is worth noting that appeals to epistemic externalism about knowledge are of no help here. By epistemic externalism about knowledge I will mean views that hold that certain contingent considerations (which do not supervene on the subject's non-factive mental states and of which the subject need have no awareness) can contribute in key ways to whether a subject has the justification or other epistemic credentials needed for knowledge. E.g., an externalist view might hold that true beliefs produced by a suitably reliable process count as knowledge. One might hold, contrary to my argument, that Ellie's belief is reliable and thus she has knowledge. The problem, however, is that a correct externalist theory must avoid attributing knowledge to subjects with certain epistemic deficiencies, like being in a Gettier case—otherwise we would simply take the cases to be counterexamples to the externalist theory. Externalist theories are appealing but notoriously difficult to formulate, and an opponent of KFF would just take the tension between the externalist theory and Ellie's epistemic deficiencies as a sign that the theory needs revised.

²⁰ See, e.g., Warfield (2005) and Fitelson (2010).

Perhaps it is just a brute fact, rather than a product of extreme confidence, that Ellie lacks all ersatz beliefs—that seems *possible*, after all.²¹ If it is right that this sort of case is the best line of defense for KFF, then the kinds of cases that KFF advocates have used don't actually establish KFF. So even if my argument against KFF fails, I would have still shown that previous arguments for KFF were unsuccessful and that the case for KFF derives from this sort of example instead.

I still think we should have misgivings, though, about TV Show 6. My argument here is more extenuated, but that is because TV Show 6 attempts to be convincing by its unspecificity, taking it as a mere brute fact that Ellie lacks all beliefs (even background or dispositional beliefs²²) that could enable an alternate evidential path. As I understand TV Show cases 2–5, Ellie does not heed a margin of error for her watch: she has extreme confidence in her watch's accuracy and forms no beliefs from it about the approximate time. So if she looked at her watch

²¹ I think that it is best to appeal to a brute fact, rather than cognitive stress. If the reason why Ellie fails to believe that it is approximately six thirty is some form of cognitive overload, then that fact might generate an epistemic defect and jeopardize any KFF intuitions about the case.

²² Here is the argument that background or dispositional beliefs can figure in K-evidential paths. Consider two cases. Case 1: say that *S* holds the background or dispositional belief that if *p* then *q*. *S* learns that *p*, and consciously reflecting on *p*, *S* infers *q*. *p* is *S*'s S-evidence, and we can stipulate that it is also *S*'s K-evidence. But clearly *p* is not *S*'s only K-evidence. The background or dispositional belief that if *p* then *q* is also an essential part of the evidential path. If *S* didn't even believe that, then clearly *S* wouldn't know that *q* (I am not using implicit quantifiers here; I am assuming a specific example in which one must know the conditional in order to have good reasons; e.g., if the orange light is lit then the heist is a go). Thus background or dispositional beliefs can be in K-evidential paths. In case 1 the background belief is on the same path as the S-evidence, but we can loosen that restriction. Case 2: say that *S* knows of two evidential paths from *p* to *r*: (i) *S* knows if *p* then *q* and if *q* then *r*, and (ii) *S* knows if *p* then *t* and if *t* then *r*. *S* then learns that *p* and, consciously reflecting on path (i), *S* infers that *r*. Path (i) is S-evidence and K-evidence for *S*, but it is also possible that path (ii) causally sustains *S*'s knowledge that *r*, and thus is a basis enabling *S*'s knowledge. If *S* came to doubt that if *p* then *q*, then *S* would still know that *r* via path (ii). If that is right, then background or dispositional beliefs on alternate evidential paths can also be K-evidence. In order for there clearly to be no alternate evidential path for Ellie's knowledge in TV Show 6, then, Ellie cannot hold even a background or dispositional belief about the approximate time.

when it said 6:59, she would still believe that her show is on—after all, she has extreme confidence in her watch’s accuracy and isn’t worried about it only being approximately correct, as normal subjects often are when they consult a time-telling device.²³ My argument against TV Show 6 then uses the following disjunction: in TV Show 6 either Ellie follows this same behavior or she doesn’t.

First, let’s say she does (I’ll call this case 6a). Then Ellie’s behavior is *akin* to being extremely confident in the accuracy of her watch, even though such confidence isn’t stipulated in the example. That is, there is still something amiss about her method of belief formation. Regardless of how the case is stipulated, either she has good reason for this behavior or she doesn’t, but either way the previous verdicts apply: if the former, she is Gettiered; if the latter, she is epistemically defective. TV Show 6a therefore makes no advance on earlier cases. Here is a further way of illustrating the point. Ellie’s method of belief formation could have easily led her to error in case 6a just like in cases 2–5, because in nearby worlds in which she looks at her watch when it reads 6:59, she will form the belief that her show is on, even though that belief is false (her watch is two minutes slow, recall, and her show ends at 7:00). That is to say, she violates safety, which I will understand as the principle that in nearby worlds in which the subject holds the belief in the same way, the belief is true.²⁴ Now, I am not claiming that safety violations are a necessary condition of every putative KFF case (much less that safety principles allow us to analyze knowledge in some significant way).²⁵ Rather I am illustrating how Ellie in TV Show 6a has some of the same epistemic problems as Ellie in cases 2–5: her method of belief

²³ Of course, in some situations she might not form the belief that her show is on because she isn’t thinking about her TV Show, but we are supposing that she is.

²⁴ This is a standard formulation, but safety principles can be formulated in many ways (for discussion see Pritchard 2007, 2009).

²⁵ I owe a debt to John Hawthorne for pressing me on related issues.

formation in 6a seems to exhibit something akin to problematic over-confidence that makes her error-prone and gives us grounds for thinking that she lacks knowledge in 6a just as she does in cases 2–5. Even if we don't stipulate in the example that Ellie has extreme confidence, if her method of belief formation is similarly problematic as it is in the previous cases, then the advocates of KFF have not become more convincing by being more elusive in how they stipulate cases.

Let's consider the other possibility, then, that Ellie in TV Show 6 does not mirror the behavior of Ellie in cases 2–5 (call this 6b). So in this case Ellie does heed a margin of error in her watch. Then I contend that our intuition that she has knowledge is guided by the reader filling in the example with an alternate evidential path, so the case is inconsistent. For consider what is necessary for Ellie to heed a margin of error. When she looks at her watch when it reads 6:59, she forms the belief that it is exactly 6:59, yet she resists the inference that her show is on (while considering whether her show is on: see footnote 23). When we have the intuition that Ellie in 6b has knowledge, that is because a subject who follows her watch like this is in good epistemic standing—here we finally have a case where there is not something epistemically amiss about the subject—but that is because we understand her behavior as manifesting the beliefs that if a watch isn't exactly right it is still approximately right and if it is merely approximately right then one shouldn't trust its extreme accuracy near certain thresholds, when the chance of error is high. The problem is that these beliefs constitute an independent veridical evidential path.²⁶ It is plausible, then, to think that the intuition that Ellie has knowledge in TV

²⁶ The path can be demonstrated this way: start with the premise that the watch reads six thirty (which Ellie uncontroversially knows). Now either it is exactly six thirty or it isn't. In the first case, clearly the show is on; in the second case, use the premise that if the watch isn't exactly right, it is still approximately right, and thus the show is on. This evidential path requires only true premises (it only uses the proposition that it is exactly 6:30 in the assumption of a

Show 6b relies on us filling in assumptions about her based on normal subjects, assumptions that are inconsistent with the case. Otherwise, if we genuinely have to hold our assumptions in abeyance, the force of an argument from intuition or appeal to cases simply evaporates, since we are not being allowed enough information to make judgments about the case. I therefore find the appeal to TV Show 6 unconvincing: either the case inherits the same problematic features of previous cases, or it violates our ability to understand the case, thereby compromising its own force.

My assessment of part 2 of the case for essential KFF, then, is this. If there is no rational epistemic explanation for why Ellie lacks beliefs that would enable an alternate evidential path, then

- (i) The mismatch between her confidence and her watch's accuracy shows, to the extent that we can make clear judgments about the case, that the epistemic features of the case suggest that she doesn't have knowledge, and
- (ii) Ellie is not as normal of a subject as KFF advocates have supposed, and our intuitions about her are, to that extent, less valuable for intuitive-based arguments for KFF.

subproof—that cannot constitute relying on a falsehood; think about argument by reductio). I should also point out that subjects don't have to go through this sort of reasoning consciously. Arnold (2013: 169) argues that in examples like TV Show 5 or 6a Ellie's true belief that the watch reads six thirty cannot enable a path independent of the falsehood that it is (exactly) six thirty. I do not need to dispute that claim, as I think that Ellie lacks knowledge in TV Show 5 and 6a. I am merely arguing that if Ellie in TV Show 6b held the belief that if the watch isn't exactly right then it is still approximately right, etc., then the watch belief could enable an independent path containing just truths.

It is therefore plausible to reject essential KFF in part 2, a conclusion which mirrors that of part

1. Contrary to the claims made by KFF advocates, we are under no compulsion to accept essential KFF.²⁷

3.2 *Non-Essential KFF*

If my argument so far is correct, then there is no essential KFF. That is already an important advance, because all KFF advocates claim that there is essential KFF.²⁸ I will now argue, furthermore, for the conditional that *if* there is no essential KFF *then* there is no KFF at all. Non-essential KFF cases would involve two or more independent K-evidential paths. Focus for simplicity on a case with two evidential paths, one of which relies on a falsehood and one of which does not. In order to be independent, the paths must not rely on each other; they must be capable of enabling the subject's knowledge alone, without the other path's support (or any other path's support).²⁹ There is therefore a non-trivial relationship between being independent and being potentially essential, since if an evidential path does not rely on any other path, it could enable the subject's knowledge alone. I will put the idea this way:

²⁷ Here is one further objection to my case against essential KFF. The KFF advocate might maintain that in TV Show 1 there is one super-K-evidential path that involves both the falsehood as well as some truths that are capable of enabling knowledge (e.g., the truth that it is approximately six thirty). The subject would intuitively have knowledge in the case, yet the super-K-evidential path involves a falsehood; thus there is KFF. Since that is the sole evidential path, it is essential KFF. This objection succumbs to the kinds of points already made: contrary to stipulation, the falsehood does not seem to be on the path which actually enables the subject's knowledge, as there is a sub-part of the super path which alone would enable knowledge and which does not depend on the falsehood. The opponent of KFF can simply maintain that our intuitions are guided by the veridical path. If that were the only possible type of essential KFF, then it is even harder to motivate the thought that the falsehood is on the K-evidential path; the fact that no other falsehood can enable knowledge corroborates the assessment that this falsehood isn't actually contributing to the knowledge.

²⁸ Coffman's (2008) view is discussed in Section 5.

²⁹ The case in which they are not independent was considered in footnote 27.

Link Principle: If K-evidential path G is independent of K-evidential path H (and any other evidential path) for subject S 's knowledge that p , then there is a possible case in which S knows that p via G alone.

Link says that being independent requires being potentially essential, because if path G cannot enable the knowledge alone, without H , then it isn't actually independent of H . If that is right then there is no KFF at all, because there could be non-essential KFF only if there could also be essential KFF; but there is no essential KFF (by hypothesis).

The only way to counter this argument is to reject Link, but that is very hard to do. The fact that all KFF advocates argue for essential KFF is perhaps explained by Link's evident force. Potential violations of it would be cases in which two evidential paths are epistemically independent but somehow metaphysically dependent, so that there is no possible world in which the falsehood-involving path enables the knowledge without the truth-involving path also enabling the knowledge; then there would be no case of essential KFF because of the metaphysical dependence, and yet the falsehood-involving path would still somehow be epistemically independent of the all-truths path. I have two replies to this possibility. First, no KFF case ever presented meets it. So if it is required for KFF, then no one has yet demonstrated that there is KFF. Second, and more importantly, a defense of KFF in this way (using Link violations) faces a fundamental worry. Firstly, a modified Link principle could be used to show that the only possible type of KFF is metaphysically dependent KFF. Secondly, the KFF advocate would have to maintain that the falsehood-involving evidential path in a metaphysically dependent case really is a genuinely epistemically independent path. But the fact that the path

cannot exist independently of the truth-involving evidential path, combined with the fact that there is no KFF in any other case, suggests the counter-explanation that KFF is indeed impossible and the falsehood involving path is not actually epistemically independent. The inability to use Link to demonstrate that impossibility in metaphysically dependent cases would appear to be a mere artifact of the metaphysical dependence, not a sign of any special epistemological features of the case.

Burden of proof considerations are relevant here too. If there is a presumption that reasoning that critically relies on a falsehood cannot yield knowledge, that presumption applies just as much to an evidential path that is the sole evidential path as it does to an evidential path that is not the sole path but is nonetheless supposed to be a path operating independently of other evidential paths. If we have intuitions that a subject has knowledge in a case of metaphysically dependent non-essential KFF, there is always both an alternate explanation for our intuitions (since there is always a K-evidential path involving all truths) and a basis for thinking that the falsehood-involving path doesn't actually enable the knowledge. If there is no essential KFF, it is therefore very plausible to think that there is no KFF at all.

Before continuing with the next section I will address one final worry. I said that an alternate evidential path would have to meet two conditions:

- (i) the subject must believe the propositions in the path, and
- (ii) the beliefs from (i) must bear the right causal or explanatory relations (or whatever basing requires) to the subject's knowledge.

In essential KFF there is no path that contains only truths and meets these conditions. When constructing clear cases of essential KFF, I followed KFF advocates and focused for simplicity on cases that violate (i) (and, thereby, (ii) as well). One might worry, though, that perhaps there could be cases of KFF that only violate (ii). This type of case would face similar objections to those I have already made, with the complicating factor that our intuitions about the cases are liable to be confused by the need for the subject to actually hold beliefs in an alternate evidential path. We can press the point this way. Either the falsehood involving path is independent of the all-truths path or it isn't. If it isn't, then it seems that the case is inconsistent, since the all-truths path is not supposed to play any basing role.³⁰ If it is independent, then there must exist a KFF case that violates (i), that is, a case in which the subject does not even hold the ersatz beliefs. For if cases violating (ii) were the only possible sort of cases, then quite plausibly our intuitions are simply being influenced by the potential veridical path. I therefore consider this sort of case even weaker than those that have already been addressed.³¹

4. Counterfactually Dependent KFF (KFF*)

Fitelson (2010) makes two refinements to standard KFF cases. First, he generates examples in which the knowledge counterfactually depends on the falsity of the premise (which Fitelson calls KFF*).

TV Show 7: This case is like TV Show 1, but now there is a broken clock on the wall that reads six thirty, and in this case Ellie first glances at the clock. A person just moved and

³⁰ Even if it is consistent, I dealt with this sort of case in footnote 27.

³¹ There is a further way of pressing the point, which is by comparing this case to Lehrer's gypsy lawyer—see Section 5 for that discussion.

temporarily obscures the clock, however, so Ellie looks at her watch. The watch says six thirty, and again she infers, from the consideration that it is (exactly) six thirty, that her favorite TV show is on.

Now Ellie's knowledge is counterfactually dependent on the falsity of her premise, in this sense: if the premise (that it is six thirty) *hadn't been* false, then it would have been six thirty and the person would not have been obscuring the clock yet. Then she would have formed her belief that it is six thirty by the clock and, since the clock is broken, she would not have known that her show is on (she would have been Gettiered, a la Russell).

This case does not present a different challenge from the cases we have already considered, however, since TV Show 7 is relevantly like TV Show 1. It is consistent with the example that Ellie knows that it is approximately six thirty, and that that knowledge is the K-evidential path that enables her knowledge that her show is on.³² Although the counterfactual dependence is interesting, it does not interfere with the way that the knowledge is based in the actual world of the example. To make a general point: counterfactual dependence on falsehood *per se* does not necessarily have any bearing on knowledge.³³ Consider:

New Restaurant: Say I falsely believe that there is a baseball game tonight, so I depart for the ballpark. On the way I learn that there is a new restaurant in town (of which I would have otherwise remained ignorant).

³² Nor does Fitelson claim any different: he argues that *if* the KFF original examples are genuine knowledge from falsehood, *then* so is KFF* (Fitelson 2010: 666–667).

³³ Cf. Klein 2008: 26.

New Restaurant involves counterfactual dependence on falsehood: I would not know that there is a new restaurant had I not held the false belief about the ballpark. Clearly, however, this kind of counterfactual dependence does not impede my knowledge about the restaurant. There is an obvious reason why that is so: the grounds that enable my knowledge about the restaurant do not involve the falsehood. According to KFF dissenters the situation is the same in TV Show 7: the K-evidential path that enables Ellie's knowledge that her show is on involves only truths (e.g., that it is approximately six thirty); it does not involve the falsehood. Now, clearly there is a difference between New Restaurant and TV Show 7: the falsehood in TV Show 7 is much more intimately involved in the knowledge. Specifically, the falsehood is on the S-evidential path! But we have no reason to be bothered by that fact: the arguments in Section 3 have already shown that a falsehood on an S-evidential path is compatible with an alternate K-evidential path actually enabling a subject's knowledge. If that falsehood has interesting counterfactual properties, so be it; counterfactual dependences are cheap—they do not necessarily interfere with K-evidential paths.

Fitelson makes a second refinement, a specific version of KFF*, which is supposed to further strengthen the case for KFF.

TV Show 8: This case is like TV Show 7; the only difference is that now Ellie is convinced of the extreme accuracy of her watch and only forms very specific beliefs about the time from it (she is convinced of the approximate accuracy of the clock on the wall).

Notice the difference: if in any nearby world Ellie were to form the belief that it is approximately six thirty, she would do so because of the *clock* (she is disposed to only form exact beliefs from her watch), and thus if her inferential belief were based on that truth it would not be knowledge (it would be Gettiered again). TV Show 8 thus blocks the alternate evidential route that I appealed to in TV Show 7.

But TV Show 8 does not improve the case for KFF, because it is relevantly like TV Show 2, where Ellie will not form approximate beliefs based on the device that is supposed to be actually enabling her knowledge. TV Show 8 therefore faces the exact same disjunction that TV Show 2 did: either there is a good epistemic reason why Ellie will only form exact beliefs about the time from her watch or there isn't. Either way, once we properly understand the case we lose the intuition that she has knowledge. Again, what matters is whether the supposed basis of her knowledge is capable of enabling knowledge in the actual world of the example; the role of the clock is a distractor that potentially interferes with our intuitions, but the clock, and the interesting counterfactual dependencies it generates, does not ultimately matter. KFF* cases therefore pose no problems other than those we have already dealt with.³⁴

5. The Ersatz Solution

My argument against KFF requires, for any putative KFF case in which the subject possesses knowledge, that there is an alternate K-evidential path containing only truths that

³⁴ Objection: in TV Show 8 the falsehood is counterfactually and *explanatorily* relevant to the subject's knowledge (Fitelson 2010: 668–669; Arnold 2013: 171), which would interfere with the ability of an alternate path of all truths to meet the basing requirements. This objection is misplaced, because I have argued that Ellie in TV Show 8 lacks knowledge; if she doesn't have knowledge then *a fortiori* she doesn't have knowledge that is explanatorily dependent in a critical way on falsehood. I said that in TV Show 7 Ellie has knowledge, but by the arguments of Section 3 that is because there is a way of understanding the case on which the veridical path meets the basing requirements. See Section 5 for an elaboration of this point.

actually enables the subject's knowledge. I will call this the Ersatz Solution (ES) to KFF, because in any putative KFF case there will be a substitute true belief whose content is the actual (and sole³⁵) K-evidence. I will first clarify how ES is different from other positions that have been advanced in the literature, and then I will defend ES from objections.

First, ES should be distinguished from Coffman's view of KFF (Coffman 2008: 190–191). Roughly, Coffman claimed that in cases of KFF the falsehood is not essential to the knowledge in the following sense: the subject holds or is disposed to hold a true belief (such as the belief that it is approximately six thirty) that could enable the subject's knowledge. The false belief isn't "essential" because there will always be a true belief like the approximate-time belief that could have led to knowledge (e.g., if Ellie had based her belief that the show is on on it instead of on the exact-time belief). But, in contrast to ES, Coffman allows, firstly, that the subject doesn't actually hold the approximate-time belief; rather, the subject might merely be disposed to form the belief, which I distinguish from holding a dispositional belief.³⁶ Secondly, even if the subject does hold the approximate-time belief, Coffman allows that in the actual world Ellie's knowledge is in fact based on the falsehood; thus she has genuine, but "non-essential", KFF.³⁷ On ES, by contrast, in the actual world Ellie's knowledge is based on the path involving only truths; if that is right there is no KFF at all, "essential" in Coffman's sense or otherwise.³⁸ ES is also distinct from a position described by Klein (2008: 30, 40–47), the

³⁵ Along with any other beliefs on the same path. It is consistent with ES that there are multiple veridical K-evidential paths, but for simplicity I am focusing on cases with two potential paths, one with a falsehood and one with no falsehood.

³⁶ Perhaps Coffman thinks about dispositional beliefs differently, so I register a qualification of this point.

³⁷ Note that I use "essential" differently than Coffman; if the falsehood is on the only K-evidential path that actually enables the knowledge then it would still be essential in my sense.

³⁸ These differences make Coffman's view potentially vulnerable to Fitelson's KFF* in a way that my view is not. Coffman does not seem to hold, as I do, that the knowledge must be based

position that in the actual world of the example the falsehood and the truth are evidentially and causally redundant for Ellie, such that her knowledge is over determined. That redundancy position also admits a type of KFF (non-essential, in my terminology), since it admits that in the actual world the falsehood does enable her knowledge, and thus is K-evidence for her, albeit redundantly so. ES does not admit KFF at all, even of this redundant form.

There are two basic worries, however, that opponents will have about ES. First, I have granted that K-evidential paths must meet certain causal or explanatory basing conditions, but what guarantee do we have that the beliefs in the propositions on the evidential path containing only truths actually meet those conditions? Second, how do we know, as ES requires, that the path containing truths is the *sole* K-evidential path? I will call the first worry the Existence Problem (since it concerns the existence of a K-evidential path containing only truths that meets the basing conditions), and the second worry the Uniqueness Problem (since it requires that the veridical K-evidential path is the sole K-evidential path, to the exclusion of the falsehood involving path).

My response to these problems is based on the arguments from Sections 3 and 4: those arguments show that once we specify the relevant facts about the case, ES will deliver the correct verdict. First take the Existence Problem. Consider an arbitrary case in which subject *S* makes an inference from false S-evidence. If we stipulate that there is no veridical K-evidential path for *S* (i.e., we stipulate that the conditions for knowledge required by ES are not met), then I have shown that there will always be epistemic defects in the case that suggest that the subject lacks knowledge (as in TV Show cases 2–5, 6a and 8). If we do not make that stipulation, though, and

on a truth, and thus if in a case like TV Show 2 (or Fitelson's KFF* analogue, TV Show 8) the knowledge is not based on a truth, it is not clear whether and why, on Coffman's view, the subject must lack knowledge. But I will not try to settle the issue.

we have the intuition that the subject has knowledge (as in TV Show cases 1 and 7), then ES only requires that the details of the under-described example are filled in in a way consistent with that knowledge.³⁹ What matters is whether my previous arguments against KFF are correct; if they are, then there is no Existence Problem for ES.

The KFF advocate might find my response dissatisfying. For example, what about the case discussed at the end of Section 3.2? In that case there are two evidential paths, one veridical and one not, and the falsehood-involving path is the S-evidence. What guarantee do we have that the veridical path meets the basing requirements (the Existence Problem)? Isn't it possible that it doesn't meet them, and yet the subject still has knowledge? My answer is no, it isn't possible. If my previous arguments are correct, then a falsehood involving path cannot enable knowledge—it is, basically, a Gettiered path. Once we understand that, the case under discussion is like the following:

Deviant Max: Max sees a dog statue that he takes to be a dog in the yard; he infers that there is an animal in the yard; there is also a squirrel in the yard; so far this case is like the previous Max case, but now we add that Max *does* know about the squirrel, though we stipulate that he doesn't (for some reason) use that knowledge to ground his belief that there is an animal in the yard.

We should judge that Deviant Max does not have knowledge—his only basis for belief, by stipulation, is flawed. He is in an unusual type of Gettier case and merely has justified true belief. The putative KFF case under discussion is the same. If we require that the veridical path

³⁹ For example, when we have the intuition that Ellie has knowledge in TV Show 1, all that is required is that the veridical K-evidential path is a sustaining cause of her belief.

doesn't meet the basing requirements, then Ellie has a justified true belief that her show is on, but she doesn't know it because, like Deviant Max, by stipulation she won't take the path that could enable her knowledge. If there is anything bizarre about these cases, it is that the subject could fail to let his or her belief be sustained (based) on all his or her sources of information. These cases are related to Lehrer's famous gypsy lawyer case (Lehrer 1990: 169, as well as many other places).⁴⁰ Many people have the intuition that the gypsy lawyer's belief that his client is innocent is not justified.⁴¹ The lesson is that even if one possesses a good evidential path, if one fails to allow it to be a basis (causal, e.g., on this intuition) then it does not pass on its epistemic credentials to one's inferred belief. The exact same goes for Deviant Max and for Ellie in the case at issue.

The Uniqueness Problem fares similarly: if my previous arguments are correct, then there is no Uniqueness Problem. The Uniqueness Problem concerns cases in which two evidential paths (one veridical, one not) meet the basing requirements, and questions how we know that the falsehood-involving path is not K-evidence. The arguments from Sections 3 and 4 aver that

⁴⁰ Summary: a gypsy lawyer believes that his client is innocent of one of eight murders because of his (the gypsy lawyer's) tarot cards. The lawyer then becomes aware of actually good evidence showing that the client is innocent of that murder, but because of his desire to find a culprit for all eight murders the actually good evidence is unable to motivate him; only his unwavering faith in the tarot cards is. Lehrer says that the actually good evidence is therefore not a cause of the lawyer's belief. Yet his belief is intuitively justified, according to Lehrer, and if it is justified our intuitions must be guided by the thought that the actually good evidence is his reason (since if just the tarot cards were his reason we would not think his belief is justified). Thus, according to Lehrer, an evidential path can be a basis without being a cause. If Lehrer is correct, then my case against KFF is even stronger, because then the veridical path *would be* a basis for Ellie's belief even if it weren't a cause (it would meet analogous conditions to the gypsy lawyer with simple elaborations) and would thus explain the supposed intuition that she has knowledge. I think Lehrer is wrong, however, which is why I don't find it plausible that Deviant Max or Ellie in this case has knowledge. As I explain below, whether Lehrer is right or wrong does not matter for the fate of ES.

⁴¹ See Audi (1983: 223–225), Goldman (1979: 22), Harman (1970: 841–855), Lemke (1986), and Swain (1981: Ch. 3).

falsehood-involving paths cannot enable knowledge. If that is right, we have reason to believe that any falsehood-involving path is not K-evidence for the subject. In a case in which a subject has knowledge and two independent evidential paths meet the basing requirements (one path with a falsehood, one without), it follows that only the path of all truths is K-evidence. An example would be the Havit-Nogot case from Section 2, in which the subject infers that someone owns a Ford from two considerations, the consideration that Havit owns a Ford (which is true) and the consideration that Nogot owns a Ford (which is false). Since falsehoods cannot enable knowledge, we know that only the Havit path is K-evidence.

Perhaps the following lemma will make the point clearer. I will argue that ES can allow⁴² that the falsehood-involving path also meets epistemic basing requirements for a subject's belief and provide her belief with positive epistemic status without that path thereby being K-evidence for her. That is, it is consistent with ES that the falsehood-involving path provides some positive epistemic status to a subject's beliefs. Define **J-evidence** to be evidence that provides some positive justification for a subject's inferential beliefs. An evidential path might meet the basing requirement and be J-evidence for a subject without thereby being K-evidence for her. When a subject has multiple sources of information, those different sources might play sustaining roles for her beliefs even if only one of the sources can enable knowledge. For example, in the Havit-Nogot case, the falsehood that Nogot owns a Ford might be J-evidence for the subject's belief that someone owns a Ford—it is a genuine evidential path that justifies the belief. But, critically, it is a Gettiered path, and it is therefore incapable of sustaining knowledge, even though it is capable of providing some epistemic support. For another example, consider a lottery case. Say a subject has a justified belief that she didn't win the

⁴² Though it isn't committed to this possibility.

lottery just based on the odds; she then comes to know that she didn't win by testimony. The information about the odds is not K-evidence for her, even in part; it is not on the evidential pathway that is the basis of her knowledge.⁴³ But it is still J-evidence for her; it still provides her belief with positive epistemic status and still contributes to the sustaining of her belief, since if she heard that the testimony was wrong she would still believe that she didn't win because of the odds.

That the falsehood-involving path could be considered J-evidence for her might further explain why philosophers have thought that there is KFF. One might hold that Ellie's false S-evidence is playing a positive epistemic role—it is J-evidence for her (or on a path of J-evidence for her). But it is like the Gettiered-Nogot path or the odds-based lottery path: it is a path that cannot enable knowledge. Rejecting KFF therefore does not require us to view Ellie's S-evidence as epistemically irrelevant; its role is just overshadowed by the K-evidence. But just as being S-evidence is not sufficient for being K-evidence, being J-evidence is not sufficient for being K-evidence (as the lottery example shows).⁴⁴

It is worth considering two final objections. One is raised by Arnold, who argues that anti-KFF accounts such as ES incur the burden of providing a positive explanation of how epistemic basing actually works in putative KFF and related cases (Arnold 2013: 171). In a trivial sense that is true. ES owes us an explanation of what is going on in the various cases; that is the sort of explanation I have already given. In another sense (which I do not necessarily

⁴³ Presuming a standard understanding of lottery cases; readers who reject that understanding can substitute a different example. For discussion see Hawthorne (2004).

⁴⁴ A different type of uniqueness problem is raised by Klein (2008: 40): perhaps there are too many truths that contend to be on the K-evidential path. This problem is simply solved by the facts, viz., the facts about which beliefs are playing sustaining roles. If many such beliefs are doing so, then there are many K-evidential paths, just as in any case of knowledge that is sustained in many ways.

attribute to Arnold), though, ES does not face a special burden here: rejecting KFF, as ES does, does not require deciding which theory of epistemic basing is correct (whether basing is causal, counterfactual, doxastic, hybrid, etc.).⁴⁵ That issue is by and large orthogonal to KFF, and neither KFF views nor anti-KFF views should be hostage to it. As my previous remarks indicate, no matter how the basing debate is settled ES will deliver the correct verdict about cases. ES is a theory about whether subjects have knowledge given facts about evidential paths (including facts about basing); it is thus downstream from the basing debate. Once we specify the relevant facts about cases, and refract those facts through the correct theory of basing, then ES issues a verdict. If the arguments that I have already given are correct, then that verdict will be right regardless of how basing was decided.

Perhaps this point will be clearer with an example. One of the major issues raised in the basing debate concerns Lehrer's gypsy lawyer, but my arguments against KFF work regardless of how that issue is settled. If one shares Lehrer's intuition that the gypsy lawyer has a justified belief that his client is innocent, then the non-causal good evidence is a basis. If that is right, then the Existence Problem for ES is made even less pressing, because as long as the subject holds a veridical ersatz belief we don't have to worry about it being causal in order for it to ground the subject's knowledge. If one doesn't share Lehrer's view, that is fine for ES too: then one simply shouldn't think that Deviant Max, or Ellie in the related cases, has knowledge. However the basing issue is settled, all that I require is that one's judgments consistently transfer to putative KFF cases.⁴⁶ The burden of proof is therefore not on ES to solve basing. Rather, the burden of proof is on the shoulders of the KFF advocate to show why one shouldn't consistently

⁴⁵ See Korcz (1997, 2010). I have been granting that basing involves some causal or explanatory condition, but none of my arguments actually depends on that.

⁴⁶ If one doesn't have consistent intuitions across these cases, then the problem is with one's intuitions, I would say, not ES. One is failing to understand the epistemic similarities.

apply one's intuitions across the cases. I have not shown that it is impossible for basing considerations to be used to argue for KFF; rather, I have shown that we have no reason to think that they can be so used and thus ES is not hostage to solving the basing debate.

The second and final worry one might have is that the points I make concerning TV Show cases, while successful, do not generalize to other putative KFF cases.⁴⁷ For example, consider a type of case from the literature that is different from Ellie's:

Christmas Morning On several successive years a child finds presents under the tree on Christmas morning and believes (from a reliable authority) that Santa brought them. She then comes to believe that this year there will be presents under the tree on Christmas morning (because, she reasons, Santa brought presents last year, and the year before that, etc.). (Cf. Klein 2008)

It is natural to think that the child knows that there will be presents under the tree, yet her S-evidence (and her K-evidence, according to the KFF advocate) is false. Since this case is rather different than TV Show, one might wonder how my account handles it. One important issue in Christmas Morning is whether the fact that the subject is a child is supposed to make an epistemic difference; whether it does or not, however, I think that ES delivers the correct verdict. First, say that the fact that the subject is a child is not at issue. Then it is natural to suppose that, even though the subject is a child, she still holds the background belief (and knows) that someone brought the presents. After all, consider what would naturally happen if her friend convincingly informed her (and she thereby came to know) that Santa isn't real. She wouldn't

⁴⁷ Thanks to an anonymous reviewer for pressing this point.

drop her belief that there will be presents under the tree; she would just drop her belief that it was Santa who was bringing them. So there is a normal way to understand the example, and thereby underwrite our intuitions, on which the subject holds true background beliefs that play a sustaining role of the subject's knowledge (as ES requires).⁴⁸ Compare the analogous point for Ellie: in TV Show 1, Ellie knows that it is approximately six thirty, and if someone told her that it wasn't exactly six thirty, she wouldn't drop her belief that her TV show is on. So if the subject in Christmas Morning is cognitively normal (that is, roughly as capable as, even though less knowledgeable than, a normal adult), then this sort of case is not fundamentally different from the TV Show cases, and thus there is no reason to think that the arguments concerning Ellie fail to generalize.⁴⁹

Next, say that the fact that the subject is a child is supposed to make an epistemic difference (for example, the child is not capable of forming the belief (or making the relevant inferences) that if Santa didn't bring the presents then at least someone did). Then I think that the KFF opponent should simply deny that we have clear and theoretically significant intuitions about a cognitively deficient subject like this genuinely having knowledge in this case (as distinct from proto-knowledge, or some other related cognitive state). The fact that there is no

⁴⁸ I think that part of the force of the example derives from this natural understanding of the case, as we all know that this happens countless times each year.

⁴⁹ For example, we could stipulate in the Christmas Morning case that there is no veridical ersatz path, but then the same points made concerning TV Show cases would apply. E.g., if she has good epistemic reason for avoiding an ersatz path, she will be Gettiered. Perhaps the strongest way of pushing the objection is to maintain (as in TV Show 6) that it is just a brute fact that the subject holds no beliefs in an ersatz path. The same considerations I marshaled against TV Show 6, however, would then apply: if, upon learning that there is no Santa, the subject would not still believe that there will be presents, then her epistemic behavior is analogous to problematic overconfidence and it is plausible to maintain that she lacks knowledge; if she would still believe that there will be presents, then our intuitions are being guided by us treating her as a normal subject who has background beliefs (see the previous footnote). Either way it is plausible to maintain that there is no KFF in the case.

KFF for normal subjects suggests that there is no KFF at all, and any intuitions to the contrary, it seems, are affected by the abnormal nature of the case. Given the natural worry that our intuitions about the cognitively deficient subject are compromised, I simply don't think that the case for KFF can rest on this sort of example.⁵⁰

6. Conclusion

I have argued that there is no KFF, essential or otherwise. Insofar as we can make clear judgments about putative KFF cases, facts about the cases suggest that the subjects actually lack knowledge, contrary to the claims of KFF advocates; to the extent that clear judgments are not possible, then to that extent the case for KFF is compromised. These arguments apply both to

⁵⁰ KFF advocates have posed other challenges as well, which they think that any anti-KFF account must face. Here are two. First, one might worry that on views like ES the knowledge in KFF cases would not be inferential, since the knowledge does not depend on the path that was actually used in the conscious inference. But intuitively the knowledge is inferential. Reply: we often consider knowledge inferential even though the premises were not consciously used in an inference. In KFF cases there is the complication that different premises were used in a conscious inference, but no reason has been given that this difference matters. Furthermore, the knowledge derives from and is based on one's other beliefs and knowledge (e.g., in TV Show 1, the belief that it is approximately six thirty); there is therefore no problem in calling it inferential. Second, Warfield (2005: 412) argued that anti-KFF accounts will have difficulty correctly adjudicating certain cases, such as Gettier cases. ES, however, has no problem here. Warfield's prime case is the original Max example: Max thinks that he sees a dog in the yard and infers that there is an animal in the yard; there is actually no dog there, but, unbeknownst to him, there is a squirrel in the yard; thus his belief is luckily true. This case is clearly not knowledge, as ES correctly entails. For it is clear that there is no substitute belief that the subject holds that enables any knowledge that there is an animal in the yard. (We are supposed to imagine here that the subject knows of no other animals in the yard—the dog appearance was all that he was going on.) Since there is no alternate K-evidential path, there is not knowledge according to ES. Similar remarks apply to any Gettier case. Warfield focused on anti-KFF views that used the following diagnosis of KFF: in every case there is a truth entailed by the falsehood that could (in suitable circumstances) enable the subject's knowledge. In his Gettier case, there is such a truth, namely the truth that there is a dog or a squirrel in the yard (which in suitable, though different, circumstances could have enabled knowledge). His case might create a problem for that diagnosis, which I have no need to dispute. But Warfield also suggests that any anti-KFF view will have similar problems adjudicating cases, and that is incorrect. ES has no such problems.

the original KFF cases and KFF* cases. Taken together they show that we should reject KFF rather than endorse it.⁵¹

References

- Arnold, A. 2013. Some Evidence is False. *Australasian Journal of Philosophy* 91: 165–172.
- Audi, R. 1983. The Causal Structure of Indirect Justification. *The Journal of Philosophy* 80: 398–415. Reprint in *The Structure of Justification*. Cambridge: Cambridge University Press, 1993.
- Bengson, J. 2013. Experimental Attacks on Intuitions and Answers. *Philosophy and Phenomenological Research* 86: 495–532.
- Cappelen, H. 2012. *Philosophy without Intuitions*. Oxford: Oxford University Press.
- Coffman, E. J. 2008. Warrant Without Truth? *Synthese* 162: 173–194.
- Fitelson, B. 2010. Strengthening the Case for Knowledge from Falsehood. *Analysis* 70: 666–669.
- Gettier, E. 1963. Is Justified True Belief Knowledge? *Analysis* 23: 121–123.
- Goldman, A. 1967. A Causal Theory of Knowing. *Journal of Philosophy* 64: 357–372.
- 1979. What is Justified Belief? In *Justification and Knowledge*, ed. G. Pappas. Dordrecht, Holland: D. Reidel.
- Harman, G. 1970. Knowledge, Reasons, and Causes. *The Journal of Philosophy* 67: 841–855.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hilpinen, R. 1971. Knowledge and Justification. *Ajatus* 33: 7–39.
- 1988. Knowledge and Conditionals. *Philosophical Perspectives* 2: 157–182.
- Klein, P. 1996. Warrant, Proper Function, Reliabilism, and Defeasibility. In *Warrant in Contemporary Epistemology*, ed. J. Kvanvig. Lanham: Rowman and Littlefield.
- 2003. How a Pyrrhonian Skeptic Might Respond to Academic Skepticism. In *The Sceptics: Contemporary Essays*, ed. S. Luper. Aldershot: Ashgate Press.
- 2008. Useful False Beliefs. In *Epistemology: New Essays*, ed. Q. Smith. Oxford: Oxford University Press.
- Korcz, K. A. 1997. Recent Work on the Basing Relation. *American Philosophical Quarterly* 34: 171–191.
- 2010. The Epistemic Basing Relation. *Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*, ed. E. N. Zalta, URL = <<http://plato.stanford.edu/archives/spr2010/entries/basing-epistemic/>>.
- Kvanvig, J. 1985. Swain On The Basing Relation. *Analysis* 45: 153–158.
- Lehrer, K. 1990. *Theory of Knowledge*. Boulder: Westview Press.
- Lemke, L. 1986. Kvanvig and Swain on the Basing Relation. *Analysis* 46: 138–144.
- Littlejohn, C. 2013. No Evidence is False. *Acta Analytica* 28: 145–159.
- Luzzi, F. 2010. Counter-Closure. *Australasian Journal of Philosophy* 88: 673–683.
- Montminy, M. Manuscript. Knowledge Despite Falsehood.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.

⁵¹ I owe a great debt to Branden Fitelson for numerous rounds of comments on this work. I also owe thanks to Martin Montminy, Peter Murphy, an incredibly helpful anonymous reviewer, and audiences at the 2014 Central APA and Montana State University.

- . 2007. Anti-Luck Epistemology. *Synthese* 158: 277–298.
- . 2009. Safety-Based Epistemology: Whither Now? *Journal of Philosophical Research* 34: 33–45.
- Pust, J. 2012. Intuition. *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), ed. E. N. Zalta, (URL = <<http://plato.stanford.edu/archives/win2012/entries/intuition/>>).
- Rizzieri, A. 2011. Evidence Does not Equal Knowledge. *Philosophical Studies* 153: 235–242.
- Sainsbury, M. 1996. Vagueness, Ignorance and Margin of Error. *British Journal for the Philosophy of Science* 46: 589–601.
- Schwitzgebel, E. 2011. Belief. *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), E. N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2011/entries/belief/>>.
- Swain, M. 1981. *Reasons and Knowledge*. Ithaca: Cornell University Press.
- Travis, C. 1997. Pragmatics. In *A Companion to the Philosophy of Language*, eds., B. Hale and C. Wright. Oxford: Basil Blackwell.
- Warfield, T. 2005. Knowledge from Falsehood. *Philosophical Perspectives* 19: 405–416.
- Weatherson, B. 2004. Luminous Margins. *Australasian Journal of Philosophy* 83: 373–383.
- Weinberg, J., Stich, S. and Nichols, S. 2001. Normativity and Epistemic Intuitions. *Philosophical Topics* 29: 429–460.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- . 2007. *The Philosophy of Philosophy*. New York: Routledge.
- Zagzebski, L. 1994. The Inescapability of Gettier Problems. *Philosophical Quarterly* 44: 65–73.