LEARNING FROM LEARNING FROM OUR MISTAKES
Clayton Littlejohn
King's College London

What can we learn from cases of knowledge from falsehood? Critics of knowledge-first epistemology have argued that these cases provide us with good reason for rejecting the knowledge accounts of evidence, justification, and the norm of belief. I shall offer a limited defense of the knowledge-first approach to these matters. Knowledge from falsehood cases should undermine our confidence in like-from-like reasoning in epistemology. Just as we should be open to the idea that knowledge can come from non-knowledge, we should be open to the idea that justified beliefs can come from unjustified beliefs.

Philosophers have long had a soft spot for like-from-like reasoning. Whatever produces a good person must be good. Whenever something is heated it is heated by something hot. The degree of perfection contained in the effect cannot exceed the degree contained in the cause. And so on. We know that there is an intuition that underlies transmission theories of causation, but we also know that the intuition is unreliable. Gin is colourless but it makes you see rainbows. Populations can become increasingly fit over time. And so on. If we were surprised to discover that knowledge can come from mistaken belief, maybe we shouldn't have been. Without some specific reason for thinking that only knowledge can beget knowledge, we should have been open to the possibility of knowledge from falsehood (KFF).

In KFF cases, a subject acquires knowledge by reasoning through a falsehood. Most of the literature on knowledge from falsehood is concerned with the *possibility* of KFF cases.[1] Some of it focuses on the significance of such cases.[2] This is a paper about the significance of the possibility.[3] For various reasons, people think that KFF cases cause trouble for the knowledge-first approach to evidence, justification, and the norm of belief. These arguments all seem to assume that certain like-from-like reasoning fails for knowledge but holds for justification. I think this is a mistake. If it's unreliable for knowledge, it's unreliable for justification.[4]

One kind of like-from-like reasoning in epistemology is reasoning that assumes that knowledge is counter-closed:

K-Counter-Closure (KCC): Necessarily, if (i) S knows that p entails q and (ii) S comes to believe q solely on the basis

---

[1] See Coffman (2008), Fitelson (2010), Klein (2008), and Warfield (2005) for defenses of the possibility. See Schnee (2015) for skepticism.
[2] See Arnold (2013), Littlejohn (2013), and Luzzi (2010).
[3] Comesana and Kantin (2010) and Rizzieri (2011) use cases of reasoning from falsehoods to try to cause trouble for E=K. Their arguments rest on assumptions similar to those at play in this paper, but they do not depend upon the intuitions that suggest that KFF is possible.
[4] I discovered that Murphy (forthcoming) is also critical of this idea. To my knowledge, he is the first person to argue in print that unjustified beliefs can serve as the basis for some inferentially justified beliefs.

of competently deducing it from p, and (iii) S knows q,
then S knows p (Luzzi 2010).[5]

Luzzi has argued rather persuasively that we can use KFF cases to show that KCC is mistaken. Let's take this as our starting point. Using these cases, some epistemologists have argued that the cases that undermine KCC also undermine various claims associated with the knowledge-first movement. They have been taken to show that knowledge isn't evidence, it isn't justified belief, and it isn't the norm of belief:

> E=K One's evidence includes p iff one knows p (Hyman 2006, Williamson 2000).
> KN: One shouldn't believe *p* unless one knows *p* (Littlejohn, 2013, Williamson 2000, Sutton 2007).
> J=K: One's belief about p is justified iff one knows p (Sutton 2007).[6]

The basic argumentative strategy is something like this. In the KFF cases we will have a pair of beliefs, one that is false and one that constitutes knowledge. The latter will be inferred from the former. For the former to give rise to knowledge, it would have to have *something* going for it. It would have to be justified, have to provide evidence, and have to conform to the norms governing belief. Since, however, the belief doesn't constitute knowledge, E=K, J=K, and KN all must be mistaken.

> The strategy only works on the assumption that justification is counter-closed:
> J-Counter-Closure (JCC): Necessarily, if (i) S knows that p entails q and (ii) S comes to justifiably q solely on the basis of competently deducing it from p, and (iii) S justifiably believes q, then S justifiably believes p.

As we will see below, the cases that cause trouble for KCC also cause trouble for JCC. If, as I shall argue below, the KFF cases give us no (new) reason to question E=K, J=K, or KN, the KFF cases are really trouble for JCC. I shall discuss the significance of this below.[7]

1.
Let's consider some representative cases from KFF literature. Three should suffice:

---

[5] Reasoning from a false premise is *not* the same thing as reasoning from a false reason. See Alvarez (2010) and Hyman (2011) for discussion of the difference.

[6] In Littlejohn (2012), I argued that doxastic justification should be understood in terms of Sosa's (2007) notion of aptness. Sosa identifies apt belief with knowledge (of a kind). Although he thought that apt beliefs didn't necessarily constitute knowledge, he's changed course. Williamson has also changed course. In his (2000), he allowed that it might be possible to justifiably believe what you don't know, but he has since embraced the view that it's impossible to justifiably believe what you don't know on the grounds that it's impossible to justifiably believe a proposition whilst violating the fundamental norm of belief. See his (2013) for discussion. He allows that there is a sense of 'justification' that is distinct from knowledge, but suggests that this sense really captures the notion of an excuse.

[7] There are, of course, familiar objections to E=K and J=K. See Coffman (2010) and Kelp (2011), for example. There have also been responses to these criticisms to consider. See Littlejohn (2012), Mitova (forthcoming), and Williamson (2013), for example. I don't want to defend E=K or J=K here. My aim is simply to argue that these cases don't give us any additional reason to think that these theses are mistaken.

Father Christmas: Virginia's parents tell her that Father Christmas will put presents under the tree for her. Believing what her parents tell her, she infers that there will be presents waiting for her tomorrow under the tree. She knows that there will be presents (Klein 2008).

Meeting Time: Alex has a 7 pm meeting, and extreme (and justified) confidence in my fancy watch's perfect accuracy. Alex makes inferences from what his watch says only if he has extreme confidence that it is perfectly accurate (perhaps he has exacting standards for what constitutes a good watch). Having lost track of the time and wanting to arrive on time for the meeting, Alex looks carefully at his watch. Because of his extreme confidence in his watch's accuracy, he reasons: 'It's exactly 2:58 pm; therefore I am not late for my 7 pm meeting'. Again he knows his conclusion, but as it happens it's exactly 2:56 pm, not 2:58 pm (Warfield 2005; Arnold 2013).

Millikan's Experiment: By measuring the electrical charges of many oil drops, Millikan determined both the value of e (i.e., the charge of the electron), and that the charges are all integer multiples of e. Determining a relatively exact value of e involved measuring the following parameters as accurately as possible: temperature, pressure, voltage, the coefficient of viscosity of air, the density of clock oil, the value of the gravitational constant and the times of rise and fall of the oil drops. The important point to note then is that all of Millikan's measurements were—and still are—taken to be evidence for the claim that electrical charge is quantized. But, the measured quantities used to determine the value of e were all approximations due to the measurement errors inherent in determining the values of the relevant parameters … So, the evidence Millikan used to confirm the claim that electric charge is quantized is only approximately true, due to the inexactness of the various methods of measurement used in the experiment (Shaffer 2013: 31. See also Hilpinen 1988: 164).

If we suppose that the protagonists have inferential knowledge as a result of reasoning from a false proposition, should we say that these cases threaten E=K or J=K?

Let's start with Shaffer's criticism of E=K. Shaffer seems to think that it's obvious that Millikan's evidence includes a falsehood, one that served as the evidential basis for Millikan's inferential knowledge. To deny this, as he sees it, is *immensely* costly:

> Adopting Williamson's or Littlejohn's truther view would then appear to commit us to the totally implausible view that the measurements made in such cases are not really evidence at all, because they are not true, and that the hypotheses that such measurements appear to support are not, in fact, confirmed on the basis of those measurements.

> This would effectively eliminate the possibility that measured values ever constitute evidence in the conduct of the science and in our more mundane epistemic lives. So, adopting the truther view of evidence with respect to propositions that report measurements undermines the very possibility of conducting rational empirical inquiry involving measurement. But, this is clearly at odds with actual practice and such measurements are evidence in both the sciences and everyday life, often exceptionally good evidence (Shaffer forthcoming: 6).

It is because Millikan's evidence included the false proposition that the value of e was $1.5924(17) \times 10^{19}$C that Millikan was able to knowingly infer that electric charge is quantized.

Klein criticizes both J=K and E=K. He thinks that the protagonists' relevant false beliefs are justified and apparently thinks that this shows that the protagonists have false evidence. In support of his claim about justificatory status, he says that this is 'an instance of the general requirement that doxastic justification which depends on another belief requires that that belief be doxastically justified' (2008: 50). On this assumption, it is thought to follow that the content of the false belief is part of the protagonists' evidence because it is a consequence of his account of doxastic justification that the content of the false belief propositionally justifies the inferential beliefs that constitute knowledge (2008: 50).[8]

2.

While I can see why someone might describe our KFF cases in ways that would cause trouble for the knowledge-first approach to evidence or justification, the grounds for thinking that this is the best way to describe such cases strikes me as rather weak.

Let's start with Millikan's Experiment. Contrary to what Shaffer thinks, I do not think it is obvious that the right thing to say about Millikan's Experiment is that Millikan came to know by reasoning from false *evidence*. On most views of evidence, the proposition p will not belong to the subject's evidence if the subject knows that p is false or should know that p is false.[9] With this very weak constraint on evidence in place, it seems that Shaffer's case is no threat. Millikan *knew*, or should have known, that the claim that the value of e was $1.5924(17) \times 10^{19}$C was nothing more than approximately true. If, as Shaffer suggests, something is false if it is merely approximately true, we should expect Millikan to have known this and known this for just the reasons that Shaffer cites. As calculating the value of e required an accurate measure of temperature, pressure, voltage, the coefficient of viscosity of air, the density of clock oil, the value of the gravitational constant and the times of rise and fall of the oil drops, it would be *mad* to think that the calculated value of e was anything more than approximately true. If it would be mad to believe p, I don't think p is part of one's evidence.

---

[8] Arnold's (2013) criticisms of J=K and E=K are similar.

[9] On the assumption that the subject knows p is false, it is plausible that the subject does not justifiably believe p. On the assumption that one should know that p is false, it is plausible that the subject does not justifiably believe p. Thus, if one thought that having a justified belief in p is necessary for having p as part of one's evidence, the proposed necessary condition on having evidence should be rather plausible.

To make the case plausible as a counterexample to E=K, we have to find a proposition that played some role in Millikan's reasoning that has these three features:

(a) The proposition was one that the protagonist could reasonably have believed;

(b) The proposition would provide adequate support for the protagonist's conclusion in that it would be reasonable for someone who believed the proposition to draw the relevant conclusion;

(c) The proposition was not one that the protagonist knew because the proposition was false.

Since the proposition that Shaffer fixes on was for Millikan either a known unknown or one that Millikan should have known couldn't be known, his proposal does not satisfy (a). The related proposition that the calculated value of e was 1.5924(17) X 10^{19}C satisfies (a) and (b), but does not satisfy (c).

In support of his description of the case, Shaffer writes:

Millikan's measurements were—and still are—taken to be evidence for the claim that electrical charge is quantized. But, the measured quantities used to determine the value of e were all approximations due to the measurement errors inherent in determining the values of the relevant parameters (2013: 31).

The expression 'Millikan's measurements' seems to me to be a bit of a fudge. It could refer to the value that was later replaced by more careful measurements or a proposition about the measurements that Millikan took (e.g., that Millikan determined the value of e to be such and such). Insofar as Millikan ran the experiment a number of times to calculate e and did not receive the same value for e in each trial, it seems like a rather odd psychological hypothesis to say that Millikan *believed* he got the right value and an odd normative claim to say that that's what he ought to have believed. Moreover, for the purposes of the experiment, getting the precise value would not matter for learning that electric charge is quantized. So long as the charges of the observed oil drops were (roughly) whole integer multiples of e, that would show that electric charge is quantized rather than continuous.

Since Millikan's Experiment does not satisfy (a)-(c), it seems not to be a terribly threatening case to E=K. And insofar as it doesn't threaten E=K, it poses no obvious threat to J=K or KN. Since, however, Father Christmas and Meeting Time do plausibly satisfy these three conditions, a further response is required.

Klein says very little in support of his contention that the protagonists' beliefs in KFF cases are justified beyond saying that there is a general requirement on doxastic justification that doxastically justified beliefs are inferentially justified only if the beliefs they are based upon are themselves doxastically justified. Thus, if he has an argument against E=K and J=K his argument rests on the undefended assumption that justification is counter-closed:

J-Counter-Closure (JCC): Necessarily, if (i) S knows that p entails q and (ii) S comes to believe q solely on the basis of competently deducing it from p, and (iii) S knows q, then S justifiably believes p.

We have to speculate a bit as to why Klein would think that JCC is true because he says little in support of it apart from the fact that he takes it to be an implication of a plausible account of doxastic justification. For my part, I would think that the discovery that KCC is false should lead us to reconsider the status of JCC. Shouldn't

the counterexamples to KCC make us sceptical of like-from-like reasoning in general? Perhaps Klein thought that we had no independent grip on what it would be for a belief to be justified apart from a belief that could justify further beliefs. This might explain some of the confidence he had in JCC. To this, I have a response. The beliefs in the useful falsehoods *do not* justify further beliefs in the ways that typical justified beliefs do. Once we see why, we can see why the cases that threaten KCC might threaten JCC.

3.
Are KFF cases any threat to E=K or J=K? Little has been said to justify the claim that they are. Perhaps it's thought that little reason is needed since we already have ample reason to think that beliefs in the useful falsehoods *could* be justified. After all, isn't there ample reason to think that there can be false, justified beliefs? If so, mustn't there be false evidence?

It would be disappointing if this were the only argument that critics of E=K or J=K could muster because then the KFF cases would not be doing much work. If we were independently convinced that J=K or E=K were false, we wouldn't care much about what we could learn from KFF cases in this regard. The interesting question is whether KFF cases put any *additional* pressure on E=K or J=K.

It might seem that Klein has provided us with at least some reason to think that KFF cases present a unique challenge to E=K and J=K. It appears that the standard account of the relationship between doxastic and propositional justification might, if combined with the idea that propositions justifiably believed are evidence (J→E), support this counter-closure principle:

> J-Counter-Closure (JCC): Necessarily, if (i) S knows that p entails q and (ii) S comes to believe q solely on the basis of competently deducing it from p, and (iii) S knows q, then S justifiably believes p.

As he sees it, doxastic justification requires propositional justification. Propositional justification, in turn, is either evidence or provided by the evidence. If we assume that the relevant beliefs are justified and accept his account of propositional and doxastic justification, we get some motivation for rejecting E=K:

> P1. If the protagonist comes to know q by reasoning from some useful falsehood p, S's belief in q is doxastically justified.
> P2. The protagonist comes to know q by reasoning from some useful falsehood p.
> C1. Thus, the protagonist's belief in q is doxastically justified.
> P3. For the protagonist's inferential belief in q to be doxastically justified, the protagonist must believe q for a good reason.
> C2. The protagonist believes q for a good reason.
> P4. The protagonist believes q for a good reason only if the protagonist believes q for the reason that p and p is part of the protagonist's evidence.
> C3. The protagonist believes q for the reason that p and p is part of the protagonist's evidence.
> P5. If p is part of the protagonist's evidence, the protagonist justifiably believes p.

C4. The protagonist justifiably believes p.

If this argument establishes (C3), E=K must be false. If the argument can establish (C4), J=K must be false. If both are false, it looks like KN must be false, too.

The thought behind (P1) is that knowledge requires doxastic justification. I'm happy to grant that this is so. In assuming that KFF cases are possible, we're assuming (P2). It would seem that (P3) and (P4) follow from the standard account of doxastic justification. The standard account says (roughly) that if the protagonist's belief is doxastically justified, it is propositionally justified by virtue of the fact that the subject has adequate reason for having that belief and the subject believes for the relevant justifying reason. The idea behind (P5) might be put this way. If a subject's evidence includes some proposition, the subject has the right to reason from that proposition. If the subject ought not reason from p, the subject must not be justified in believing p in the first place.

While this argument seems initially quite promising, defenders of E=K and J=K have a plausible line of response. Anyone who takes seriously the possibility that J=K might be true would think of proper belief or permissibly held belief as one that conforms to the knowledge norm. In other words, they identify knowing p with permissibly believing p. In discovering that KCC is false, they would reject JCC. If that requires rejecting the standard account of doxastic justification, this might well be a price they'd be willing to pay. They would reject (P3) on the assumption that the good reason that the protagonist must have for believing q is some prior reason that figured in reasoning because they'd deny that p is any reason at all by virtue of the fact that it's false.

How plausible *is* the standard account of the relationship between propositional and doxastic justification? Should we say that doxastic justification is a matter of propositional justification plus basing and that propositional justification is determined by the subject's evidence? We will see that KFF cases cause trouble for the standard story about propositional and doxastic justification, but it is also worth mentioning that there are other sources of difficulty, too. Anyone who accepts E=K will already have to say that the standard story has to be modified for cases of non-inferential belief. Consider the kinds of cases that foundationalists would use to elicit the intuitions that underwrite the isolation objection. You don't know what tonight's dish will be but when the cloche is lifted you see that they're serving fish. Prior to this experience, you had no good reason to believe that it would be fish but now you know that it is fish. If E=K is correct and the case is as I've described it, there was nothing in your possession that could have been evidence for this belief that you possessed independently from the belief. If we think of propositional justification as (a) the whatever-it-is-that-warrants-the-formation-of-a-belief where this can be possessed without the relevant belief and (b) something that supervenes upon the evidence, we would have to reject E=K if we wanted to say that in seeing the fish and coming to believe there was fish for dinner you knew and justifiably believed that it would be fish for dinner. If we accepted E=K, however, we would say that there is nothing that plays both the (a) and (b) roles. There is something that ensures that you're in a position to know that fish is being served and so there is something that ensures that you're entitled to form that belief, but that won't be evidence. So, if propositional justification is understood in terms of (a), we have to reject (b). If it is understood in terms of (b), we have to reject (a). Since we have to give up one of these to understand what happens in cases of knowledge or justified belief not acquired by using independently held evidence, we shouldn't use the standard story

about propositional justification and doxastic justification in an argument against E=K. In using it, we're already assuming that E=K is mistaken.

There are two arguments that one might offer in support of the knowledge-first views under attack. The first argument is *the ability argument*. This argument is designed to show that the KFF cases are not cases in which someone believes for a good reason because they are not cases in which the subject's inferring q is a case of believing q for a reason. The second argument is *the generality argument*. The argument is designed to show on independent grounds that the useful falsehoods considered thus far are *not* propositions justifiably believed and plausibly taken to be part of the protagonists' evidence.

3.1

The ability argument requires some set up. We often treat the subject's evidence and the subject's reasons for believing as basically interchangeable. If Alex has evidence to believe that the meeting hasn't started, he has a reason to believe that the meeting hasn't started. The reason isn't idle. It's important to the example that Alex believes that the meeting hasn't started *for* the relevant reason. This of course implies a number of things: that Alex has a certain reason, Alex exercised his capacity for identifying and responding to reasons in the course of forming the relevant belief, and that his believing what he does is a case of believing for a particular reason. The reason for which Alex is said to believe that the meeting has not started is that it is exactly 2:58. Note that the success of the argument against E=K and J=K turns on two things:

> (a) Alex believes that the meeting has not started yet for the reason that it is exactly 2:58;
> (b) This reason is not just the reason for which Alex believes that the meeting has started, it is a good reason to believe this, a piece of evidence that constitutes a justifying reason.

If one were to deny (a) or (b), one would have to reject one of the assumptions in the argument against E=K or J=K because one would have to deny that doxastic justification requires believing for a reason that's a good reason to so believe.

If Alex can reason from his belief that it is exactly 2:58 to the conclusion that the 7:00 meeting has not yet started, he has exercised a certain capacity to reason and has an ability to respond to things he believes by forming further beliefs. Does he manifest the ability to believe something for the reason that it is exactly 2:58? While one might think that he does if one thinks of reasons as propositions, representations, or propositional attitudes, this is surely not the way that a defender of E=K would think of it. They would say (or should say!) that to believe for the reason that it is exactly 2:58 one must have the overall ability to believe for *that* reason where that reason is a fact, not (just) a proposition, representation, or propositional attitude. Since we've stipulated that the relevant proposition is false, the proponent of E=K will say that Alex's belief about the meeting is not a case of believing something for the reason that it is exactly 2:58. While Alex surely has a general ability to form beliefs in light of things he believes, he surely cannot have the opportunity to respond to the fact that it is 2:58 because it is *not* 2:58.[10] Alex isn't magic. He cannot respond to facts that are not facts.

_____

[10] Kenny (1992) distinguishes general ability from overall ability on the grounds that the latter requires opportunity. A subject retains certain general abilities (e.g., the

There is a data point that critics of E=K have ignored or overlooked, one that has to do with the relationship between knowledge and the abilities involved in responding to reasons. Upon seeing that his watch reads 2:58, let's suppose that Alex comes to believe that it is 2:58, is now happy because he believes that it is 2:58, and walks to the bus stop rather than hailing a cab. If asked, he might describe his attitudes and behaviour as follows:

> (1) Why do I believe that the meeting will not start for a while? Well, for the reason that it is now 2:58. As you'll recall, the meeting doesn't start until 7:00.
> (2) I am happy that it is 2:58. That means that the meeting will not start for hours.
> (3) I shall take the bus for the reason that it is 2:58. There is no need to hail a cab. I have plenty of time to get to the office.

We can suppose that Alex is not self-deceived. He does not believe for reasons other than those he reports, does not act for reasons other than those he reports, or feel happy for reasons other than those he would mention. So, taking Alex at his word, we can restate these in the third-person:

> (1') Why does Alex believe that the meeting will not start for a while? Well, for the reason that it is now 2:58. As you'll recall, the meeting doesn't start until 7:00.
> (2') Alex is happy that it is 2:58.
> (3') Alex took the bus for the reason that it is 2:58. There is no need for him to hail a cab because he has plenty of time to get to the office.

If Alex can believe for certain reasons and these beliefs, given how Alex is otherwise, is sufficient for producing a certain affective response or a bit of behaviour, he should be able to act and feel for the very same reasons he believes. If this is right, then the argument against E=K fails. Given the details of the case, it looks like claims about the reasons for which Alex believes that specify his reason as the false proposition that it is 2:58 will be true only if (2), (2'), (3), and (3') are true. They are false.

It's easiest to see this with the case of emotion, I think. Is Alex happy that it is 2:58? He cannot be happy that it is 2:58 if it is 2:56. As Gordon (1987) and Unger (1975) have argued, Alex can only be happy that p, regret that p, be angry that p, etc. if Alex knows p. The same holds true for acting for the reason that p. Alex cannot act for the reason that p unless he knows p. The linguistic data on this seems to be rather solid. Consider:

> (2'a) Although Alex does not know that it is 2:58, he is happy that it is 2:58.
> (2'b) Alex is happy that it is 2:58. Not only that, he knows that it is 2:58.

It seems like (2'a) is false. That's some evidence that (2') is true only if Alex knows that it is 2:58. As for (2'b), it seems like a redundant conjunction. That's some evidence that the information contained in the ascription of emotion contains all of the

---

ability to reason) when she is mistaken about the facts, but we cannot say that she has the overall ability to act, believe, or affectively respond to a feature of the circumstance if she's wrong about the relevant features just like one cannot form friendships with ghosts or correctly report events that didn't take place.

information contained in the knowledge ascription. Finally, Gordon (1987) reminds us that (2') seems to be equivalent to the following:

>   (2'c) It is 2:58. That is why Alex is happy.
>   (2'd) Alex is happy because it is 2:58.

Clearly, these claims are true only if it is 2:58. If they are indeed equivalent to (2') and (2), Alex cannot be happy that it is 2:58.

The same points hold true for the cases of acting and believing for reasons. When we try to say that someone acted or believed for the reason that p and then conjoin a knowledge denial, we get the same problematic results we get with (2'a). Of course, people realize this and quickly try to rewrite these claims so that the reasons for which the agent acts or believes will be a fact about the relevant individual's psychology rather than the situation. The case of emotion is helpful here because the psychologized redescription of the subject's reasons come out as clearly false in those cases. We would not say that Alex was happy that he believed it was 2:58. Alex isn't weird, just mistaken. Two points should be stressed here. First, *if* we did want to say that Alex's reason for believing or acting was *that he believed it was 2:58*, then we *would* have to rewrite (2') accordingly. Just as we know that the fact that Alex believes such and such is not what would make him happy, we know that it would not make him act or convince him that the meeting was not going to start for a while. Second, Alex would presumably *know* these facts about his own psychology if there were any plausibility to the idea that he acted or believed for the reason that he believed something. An argument from error against E=K requires that the evidence is some proposition that Alex believes that's false.[11]

The *ability argument* can be summed up as follows. The KFF cases are counterexamples to E=K only if they are cases of believing for a reason where the relevant reason is a falsehood. For these cases to be cases of believing for a reason that's a falsehood, one has to have the ability to believe for the relevant reason. If one had this ability and having the relevant belief was sufficient for producing an affective response or would result in the agent's acting or forming an intention to act, the agent could be happy or angry that the relevant fact obtained or act for the reason that the fact obtains. One doesn't have these abilities, however. The upshot is that 'S believes something for the reason that p' entails that S knows p. While the subject might have the *general* ability to form beliefs on the basis of her belief that p when she doesn't know p, she doesn't have the *overall* ability to believe anything (or feel anything or do anything) for the reason that p because that requires the truth of p and the opportunity to discover it.

3.2

Recall that the success of Klein's argument against E=K and J=K depends upon whether the protagonist in KFF cases believes something for a reason that's a falsehood, the relevant falsehood is part of the protagonist's evidence, and the relevant falsehood is evidence for the protagonist's relevant inferential beliefs (e.g., Virginia's belief that she'll receive presents or Alex's belief that the 7:00 meeting has not started). The ability argument was supposed to show that the protagonists in the

---

[11] Does that mean that Alex does not believe or act for any reason when he acts or believes on a false belief? Alvarez (2010) and Littlejohn (2012) say that this is the right way to describe the case. Alvarez helpfully distinguishes psychological explanations of the agent's acts, attitudes, and emotions from reasons explanations. Hornsby (2007) offers a disjunctivist alternative.

KFF cases do not have the ability to believe something for a reason when that reason is a falsehood. The *generality argument*, which I shall discuss in this section, is an argument that's supposed to show that the protagonist's evidence does not include the useful falsehoods that the protagonists reason from in acquiring knowledge.

If we were to suppose, as the argument against E=K and J=K requires, that Alex and Virginia's mistaken beliefs were justified and that the propositions justifiably believed to be true were part of their evidence, it seems that we would have to deny a plausible closure principle for justification:

> J-K-Closure (JKC): Necessarily, if (i) S knows that p entails q and (ii) S justifiably believes p, (iii) S can come to justifiably believe q on the basis of competently deducing it from p.

Why is that? It seems that there are a number of obvious consequences of Alex and Virginia's beliefs that they cannot justifiably believe. Consider the following inference:

> It is exactly 2:58
> <u>If it is exactly 2:58, it is not 2:57: 58</u>.
> Thus, it is not 2:57: 58.

Whereas Alex knew that it was not 7:00, it seems rather intuitive that he does not know the conclusion of this inference. Moreover, I don't think he should reason from the premise that it is exactly 2:58 to the conclusion that it is not 2:57:58 in this case. While he's justified in believing that it is not 7:00, I don't think he is justified in believing that it is not two seconds prior to 2:58.

Consider the following inference:

> Father Christmas will deliver presents.
> If Tim says otherwise, he is either lying or doesn't know
> <u>what he's talking about</u>.
> Thus, if Tim says that Father Christmas will not deliver presents, he is either lying or doesn't know what he's talking about.

I don't think that Virginia should reason in this way *precisely* because I'd object to her reasoning from the first premise to this conclusion.

If the relevant falsehoods really were pieces of evidence justifiably believed, they should serve as a proper basis for the known consequences of those propositions. Intuitively, however, this is not what we find. What seems rather intuitive is that these subjects should not rely on the relevant false premise in inferences like the ones just described. What this suggests is that the intuition that these subjects have knowledge in Meeting Time and Father Christmas is *not* due to the fact that these subjects have a useful premise for reasoning about the world (which they don't). As an alternative conjecture, the reason that we find the reasoning acceptable is *precisely* because we think of the reasoning as a possible source of knowledge.

Speaking just for myself, the reason that I find the intuitions that underwrite treating Meeting Time and Father Christmas as counterexamples to KCC is the thought that by reasoning from the false premise, the subject's are reasoning reasonably and to a conclusion they couldn't easily have been wrong about. One reason that I think the relevant falsehoods aren't genuine pieces of evidence is that they cannot justify drawing conclusions when the subject could have easily been wrong in drawing them. This is what one would expect if the knowledge were something in the neighbourhood of a belief with a safe basis and one subscribed to a kind of knowledge-first approach to these matters. Let the critics of knowledge-first

try their hand at explaining why it should be that an alleged piece of evidence is properly included in reasoning to some known logical consequences rather than others and I expect that they will start to see the virtues of the knowledge-first approach.

The *generality argument* can be summed up as follows. The KFF cases considered are not plausible counterexamples to E=K. For these cases to be counterexamples to E=K, the falsehoods have to be part of the protagonists' evidence and generally available to them to serve as the proper basis of inference to the known logical consequences of these propositions. This, however, is not what we find. While it seems proper for Virginia and Alex to believe *some* of the known logical consequences of the relevant falsehoods, it would also be improper for them to believe some of the known consequences of the relevant falsehoods. That suggests that their evidence does *not* include the relevant falsehoods.

4.

In this section, I want to discuss the broader significance of KFF cases if, indeed, the knowledge-first theses I'm interested in (i.e., E=K, J=K, and KN) withstand scrutiny.

Let me begin by offering an argument for J=K, an argument that would show that KCC and JCC are equivalent. Suppose that one justifiably believed p where this was a case of non-inferential justification and then one infers q by means of competent deduction. Might one thereby come to justifiably believe q without knowing q? It would seem that if one subsequently forgot p, one could still justifiably believe q. One couldn't then, however, believe anything for the reason that p. JKC says that if, in this state, one knew that r followed from q, one could justifiably infer r. It seems that one can justifiably infer r only if there's some fact that one has in mind that's the reason for which one believes r. It looks like q is the natural candidate. It looks like q would have to be known.[12]

If we accept J=K, it looks like we should describe our KFF cases as follows:

> (a) The protagonist should not believe the premises that she reasons from in the course of learning that p.
> (b) The protagonist's belief in the conclusion, p, is justified.

It looks like there's a promising argument for J=K from E=K. Assuming that one can justifiably believe only if one's belief (i.e., what's believed) is itself a reason that can support further beliefs, one's belief about p is justified only if what's believed is known to be true.

If (a) and (b) are both correct, the lesson to take from KFF cases is not that there is something wrong with the knowledge-first approach. The knowledge-first approach can explain why it is sometimes acceptable to conclude that something is so on the basis of a false belief and when it would be wrong to do so. It would be wrong to do so if the inferential belief is not safe. The intuition that it is acceptable to hold the belief based on the inference from a falsehood is probably due to the realization that the inferential belief could not easily be mistaken.

Now, some people might have the following intuition. Virginia, let's say, shouldn't believe that she'll have presents for no reason at all. Moreover, her only available 'reason' seems to be a falsehood. If KFF cases are indeed counterexamples to JCC, how can we vindicate these intuitions? If, in keeping with KN, she abandons her belief that Santa will bring presents, shouldn't she also abandon her belief that she

---

[12] The crucial step here is the idea that evidence can be acquired via inductive inference. For discussion, see Bird (2004).

will receive presents? If she should abandon that belief, then either she doesn't know that she'll receive presents or her belief that she will receive presents is not justified in spite of the fact that it constitutes knowledge?

If, in keeping with KN, we identify permissible belief with knowledge and impermissible belief with belief that doesn't constitute knowledge, we can vindicate these intuitions. It is true that she should not form the inferential belief that she'll receive presents without believing that Santa will bring them. If she didn't do that, she wouldn't have come to know that she'd get presents. KN says that you shouldn't infer that something is so unless you'd therein come to know that it's so. (It's *not* true that Virginia shouldn't believe she'll receive presents unless she believes that Santa will bring them. If she knows that she'll receive presents and then forgets why she believes this, she can retain knowledge and the right to believe this.) KN vindicates the intuition that if she hadn't believed Santa would bring presents, she shouldn't believe she'd receive presents.

If KFF cases are indeed counterexamples to KCC, they would seem to show that a certain evidentialist view must be mistaken. The core commitment of Conee and Feldman's (2004) evidentialist approach to justification is that justified beliefs get that status by virtue of being supported by the subject's evidence. If the evidentialist view is combined with a factualist treatment of evidence, a view on which all evidence consists of facts or true propositions, it's easy to see that this version of the evidentialist view is in serious trouble:

> FE: Every justified belief is justified because it is supported by some piece of evidence the subject has that consists of a true proposition that is distinct from the proposition believed and part of the subject's evidence.[13]

If we dispense with JCC along with KCC, there will be cases where a subject's belief gets to be justified in spite of the fact that the subject doesn't hold that belief on the basis of further facts that support that belief. Since FE insists that the only way that a justified belief gets to be justified is by virtue of being supported by a body of evidence that consists of facts, FE must be mistaken if there are indeed counterexamples to JCC.

Of course, the evidentialists like Conee and Feldman might not be bothered by this result because they don't subscribe to FE. They don't (currently) think that a subject's evidence consists of facts or true propositions (Conee and Feldman 2011: 321-4). If the evidentialists are forced to abandon FE, it seems that they must either embrace a view that allows for 'false evidence' or a view that identifies a subject's evidence with something that isn't a proposition.[14]

There are two problems with these evidentialist views. The first problem is the obvious problem. There are powerful arguments in the literature for thinking that some evidence, at least, consists of propositions and for thinking that all such

---

[13] Cases of non-inferential justification seem to constitute a problem for this view. Even if FE is restricted to cases of inferential justification, the view seems to be problematic because there can be cases in which a false non-inferential belief is the sole basis for an inferential belief that appears to constitute knowledge.

[14] For defenses of 'false evidence' (i.e., evidence that consists of falsehoods), see Arnold (2013), Fantl and McGrath (2009), and Rizzieri (2011). For arguments that all propositional evidence must consist of truths, see Leite (2013), Littlejohn (2011), and Unger (1975). For arguments that evidence is indeed propositional, see Dougherty (2011) and Williamson (2000).

evidence must be true. These arguments suggest that evidentialists ought to embrace FE rather than try to fend off objections by embracing an objectionable treatment of what our evidence consists in. The second is less obvious. Evidentialists often present their views as if they are intuitive or as if they provide the obviously correct treatment of a case. Once we see that the evidentialist position is untenable if it is combined with an independently plausible account of what evidence consists in, we should see that the core evidentialist commitment isn't really the truism its defenders often present it as being.

If we reject the evidentialist view, we have to reject the idea that justified beliefs get to be justified by virtue of the support they receive from the subject's evidence. Sometimes that's why they're justified, but there's no essential connection between the status of being justified and the property of being supported by evidence. I'm not the first to question this picture. Externalists have long argued that it's a mistake to try to think of all justification as being grounded in relations between reasons and beliefs. Here we see that there is an independent reason to take the externalist worries seriously.[15]

By denying JCC, I'm denying the following conditional: If one shouldn't believe the premises, one shouldn't believe the conclusion. The lesson is an anti-evidentialist lesson that's in keeping with knowledge-first epistemology. Pedigree matters only to the extent that it matters to acquiring knowledge. When evidence is needed for acquiring knowledge, only knowledge will fulfil one's needs. Sometimes one can come to have a safe belief by reasoning through falsehoods and thereby acquire knowledge. This just shows that there are normatively significant relations that don't obtain by virtue of relations between beliefs and prior evidence.

The assumption that pedigree must matter to justification figures prominently in the literature on foundationalism. In discussions of the regress argument for foundationalism, for example, it's not uncommon to see the argument formulated in such a way that all justified beliefs either are non-inferentially justified or derive their justification from another justified belief. Fumerton appeals to his principle of inferential justification, for example, to generate a regress argument for foundationalism and his principle clearly requires that any justified inferential belief must derive from some further justified belief that supports it (1995: 36). If JCC is false, these options are not exhaustive. Some beliefs might be justified by virtue of support they receive from further beliefs that are not themselves justified. Thus, it looks like a common version of the regress argument for foundationalism is unsound.

Of course, the foundationalist might well admit that the argument is unsound and embrace the possibility that there are counterexamples to JCC. If anything, the counterexamples to JCC might alleviate some of the pressures that the foundationalists have to contend with. It is exceptionally difficult, for example, to say what it is about properly basic beliefs by virtue of which these beliefs to stop regresses. If there are counterexamples to JCC, the power to stop a regress has to be distinguished from the possibility of being justified on the basis of something that isn't a belief. Thus, the critics of foundationalism cannot argue that an infinite chain of reasons or some web of mutually supporting belief is necessary for justification simply by arguing against the possibility of non-inferentially justified belief. For even if there is no non-inferentially justified belief in some chain of belief, the fact that JCC is false means that it might be possible for there to be beliefs 'upstream' based on an unjustified belief that are nevertheless justified. There are interesting issues to

---

[15] See Greco (2004), for example.

be explored here if, indeed, JCC is false and the standard platitudes about pedigree prove to be mistaken.

Conclusion

The thought that KFF cases provide us with new reasons to reject knowledge-first approaches to evidence, justification, or the norm of belief assume that like-from-like reasoning fails when it's reasoning about knowledge but succeeds when it's reasoning about justification. This faith in like-from-like reasoning is unwarranted. Knowledge does beget knowledge, but it need not be begotten by knowledge. This is true for justification, too.

References:

Alvarez, Maria. 2010. *Kinds of Reasons*. Oxford University Press.

Arnold, Alex. 2013. Some Evidence is False. *Australasian Journal of Philosophy* 91: 165-72.

Bird, Alexander. 2004. Is Evidence Non-Inferential? *Philosophical Quarterly* 54: 252-65.

Coffman, E.J. 2008. Warrant without Truth? *Synthese* 162: 173–94.

Coffman, E.J. 2010. Is Justified Belief Knowledge? *Philosophical Books* 51: 1-21.

Comesaña, Juan, & Holly Kantin. (2010). Is Evidence Knowledge? *Philosophy and Phenomenological Research* 80: 447–454.

Conee, Earl and Richard Feldman. 2004. *Evidentialism*. Oxford University Press.

Conee, Earl and Richard Feldman. 2011. Replies. In T. Dougherty (ed.) *Evidentialism and its Discontents*. Oxford University Press, pp. 283-324.

Dougherty, Trent. 2011. In Defense of Propositionalism about Evidence. In T. Dougherty (ed.) *Evidentialism and its Discontents*. Oxford University Press, pp. 226-33.

Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World.* Oxford University Press.

Fitelson, Branden. 2010. Strengthening the Case for Knowledge from Falsehood. *Analysis* 70: 666-69.

Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Rowman and Littlefield.

Gordon, Robert. 1987. *The Structure of Emotions*. Cambridge University Press.

Greco, John. 2004. *Putting Skeptics in their Place.* Cambridge University Press.

Hilpinen, Risto. 1988. Knowledge and Conditionals. *Philosophical Perspectives* 2: 157-82.

Hornsby, Jennifer. 2007. A Disjunctive Conception of Acting for Reasons. In A. Haddock and F. MacPherson (ed.), *Disjunctivism: Perception, Action, and Knowledge*, pp. 244-61.

Hyman, John. 2006. Knowledge and Evidence. *Mind* 115: 891–916.

Hyman, John. 2011. Acting for Reasons: Reply to Dancy. *Frontiers of Philosophy in China* 6: 358-68.

Kelp, Christopher. 2011. Not Without Justification. *Dialectica* 65: 581-95.

Kenny, Anthony. 1992. *The Metaphysics of Mind*. Oxford University Press.

Klein, Peter. 2004. Useful Falsehoods. In Q. Smith (ed.), *New Essays in Epistemology*. Oxford University Press.

Leite, Adam. 2013. But That's not Evidence; It's not Even True! *Philosophical Quarterly* 63: 81-104.

Littlejohn, Clayton. 2011. Knowledge and Evidence. *Erkenntnis* 74: 241-62.

Littlejohn, Clayton. 2012. *Justification and the Truth-Connection.* Cambridge University Press.

Littlejohn, Clayton. 2013. No Evidence is False. *Acta Analytica* 28: 145-59.

Littlejohn, Clayton. 2013b. The Russellian Retreat. *Proceedings of the Aristotelian Society* 113: 293-320.

Luzzi, Federico. 2010. Counter-Closure. *Australasian Journal of Philosophy* 88: 673-83.

McCain, Kevin. 2014. *Evidentialism and Epistemic Justification*. Routledge.

Mantel, Susanne. Forthcoming. Acting for Reasons, Apt Action, and Knowledge. *Synthese*.

Marcus, Eric. 2012. *Rational Causation*. Harvard University Press.

Mitova, Veli. Forthcoming. Truthy Psychologism about Evidence. *Philosophical Studies*.

Murphy, Peter. Forthcoming. Justified Belief from Unjustified Belief. *Pacific Philosophical Quarterly*.

Rizzieri, Aaron. 2011. Evidence Does not equal Knowledge. *Philosophical Studies* 153: 235–242.

Schnee, Ian. 2015. There is no Knowledge from Falsehood. *Episteme* 12: 53-74.

Shaffer, Michael. 2013. E Does not Equal K. *The Reasoner* 7: 30-31.

Shaffer, Michael. Forthcoming. Approximate Truth, Quasi-Factivity, and Evidence. *Acta Analytica*.

Sutton, Jonathan. 2007. *Without Justification*. MIT University Press.

Unger, Peter. (1975). *Ignorance*. Oxford University Press.

Warfield, Ted. 2005. Knowledge from Falsehood. *Philosophical Perspectives* 19: 405–16.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.

Williamson, Timothy. 2013. Gettier Cases in Epistemic Logic. *Inquiry* 56: 1-14.