

ANOTHER BLOW TO KNOWLEDGE FROM KNOWLEDGE

Peter Murphy

ABSTRACT: A novel argument is offered against the following popular condition on inferential knowledge: a person inferentially knows a conclusion only if they know each of the claims from which they essentially inferred that conclusion. The epistemology of conditional proof reveals that we sometimes come to know conditionals by inferring them from assumptions rather than beliefs. Since knowledge requires belief, cases of knowing via conditional proof refute the popular knowledge from knowledge condition. It also suggests more radical cases against the condition and it brings to light the under-recognized category of inferential basic knowledge.

KEYWORDS: conditional proof, inference, knowledge, Peter Klein, Federico Luzzi, Ted Warfield

Alongside perceptual, testimonial, memorial, and other kinds of knowledge is inferential knowledge. According to a simple and popular view, a person inferentially knows a conclusion only if they know each of the claims from which they essentially inferred that conclusion. This is the knowledge from knowledge, or KFK, principle.¹ Intuitive though it might be, KFK has recently come under attack. This paper adds to that attack by striking another blow to KFK.

On the assumption that knowledge is true warranted belief, KFK entails that a person inferentially knows some conclusion only if each claim that they essentially inferred their conclusion from is (i) true, (ii) they believe it, and (iii) that belief is warranted.² Ted Warfield and Peter Klein have taken aim at (i) with cases

¹ For just a few of the many endorsements of KFK in the literature, see Robert Audi, *Epistemology: A Contemporary Introduction to the Theory of Knowledge* (New York: Routledge, 2003), 164, Michael Bergmann, *Justification Without Awareness* (New York: Oxford University Press, 2006), 185 and 228, Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge, Mass: Harvard University Press, 1985), 18, Alvin Goldman, "Immediate Justification and Process Reliabilism," in *Epistemology: New Essays*, ed. Quentin Smith (New York: Oxford University Press, 2008), 64, Hillary Kornblith, "Beyond Foundationalism and the Coherence Theory," *Journal of Philosophy* 72 (1980): 603.

² If knowledge is not analyzable, as Timothy Williamson argues in *Knowledge and Its Limits* (New York: Oxford University Press, 2000), KFK has to be evaluated differently. There are two possibilities here. If there is one or more necessary condition on knowledge, but knowledge is not fully analyzable, KFK can be refuted by a case in which a conclusion is known, but some

Peter Murphy

of inferential knowledge that are essentially inferred from a false belief.³ Federico Luzzi has taken aim at (iii) with a case of inferential knowledge that is essentially inferred from a belief that is unwarranted because it is Gettiered.⁴ This paper completes the medley by taking aim at (ii) with a case of inferential knowledge that is essentially inferred from a claim that is not believed.

1. The Epistemology of a Conditional Proof

Consider a simple conditional proof. A person begins by assuming some claim; then they make an inference; then they end by drawing as their conclusion, a conditional statement that features the initial assumption in its antecedent position. Here are two simple examples.

LOGIC STUDENT: A logic student is asked whether $(p \ \& \ q) \rightarrow (p \vee r)$ is true. She works her way through a conditional proof: she first assumes $p \ \& \ q$; from this, she infers p by conjunction elimination; then from p , she infers $p \vee r$ by disjunction introduction. She then concludes, and comes to believe, $(p \ \& \ q) \rightarrow (p \vee r)$.

DANA'S BIRTHDAY: Having no idea what day of the week Dana was born on, I assume that she was born on a Tuesday; I then infer that on this assumption, she was born on a weekday; then I end by concluding, and coming to believe, that if Dana was born on a Tuesday then she was born on a weekday.

While I think both of these cases disprove KFK, I will focus on the second. It refutes KFK if each of the following is true: I end up inferentially knowing the concluding conditional claim; I essentially inferred that claim from my assumption that Dana was born on a Tuesday; and I don't believe that Dana was born on a Tuesday. I devote a paragraph to supporting each of these.

All but radical skeptics can easily fill in this case so that I end up with inferential knowledge of the concluding conditional. Take the knowledge part first. Different theorists will fill in the case in different ways. Some will want to ensure that the cognitive processes that I utilized were reliable; some will want to ensure that I was in other mental states as well, like various seeming states; others will want to ensure that my belief in the conditional adheres to the truth; and yet others

necessary condition on knowledge was not met at a premise from which the person essentially inferred this conclusion. If there are no necessary conditions on knowledge, some other heuristic will have to be used to assess the person's standing with respect to each of her premises. See footnote 7 for one such heuristic and reason to think that Williamson will go along with my counterexample to KFK.

³ Ted Warfield, "Knowledge from Falsehood," *Philosophical Perspectives* 19 (2005): 405-416, Peter Klein, "Useful Falsehoods," in *Epistemology: New Essays*, ed. Quentin Smith, 25-61.

⁴ Federico Luzzi, "Counter-Closure," *Australasian Journal of Philosophy* 88 (2010): 673-683.

will want to ensure that my belief could not easily have been mistaken.⁵ All though can supplement the case in some way or other so that I end up with knowledge of the conditional. What about the further claim that I have inferential knowledge of the conditional? Since I arrive at my belief in the conditional by just one route, inference, it is known inferentially. Further support for this comes from noticing that this piece of knowledge cannot be plausibly subsumed under any other kind of knowledge: it is clearly not an instance of perceptual knowledge, testimonial knowledge, memorial knowledge, etc.

Next I essentially inferred the conditional from my assumption that Dana was born on Tuesday. Let's first see why proponents of the basic idea behind KFK should formulate their principle so that it only covers cases in which a conclusion-belief is *essentially inferred* from other claims. This requirement is needed to handle cases that involve an important kind of overdetermination. Suppose for example that my belief in the same conditional was also arrived at by another route, say someone's testimony, and suppose that each of these routes was sufficient on its own to have caused me to believe the conditional. In this variant on Dana's Birthday, my belief in the conditional would not have been essentially inferred from my assumption that Dana was born on a Tuesday. Proponents of KFK can allow that in this case, I would know the conditional even though I inferred it from one, or more, failures to know since they can allow that the testimonial route was sufficient to yield knowledge.⁶ None of this is germane to our original case though since that case does not involve this kind of overdetermination.

Last it is obvious that I do not believe, but instead merely assume that Dana was born on a Tuesday. It is easy to fill out the case by imagining that I have no evidence about what day of the week Dana was born on, and that I suspend judgment about this rather than form a belief about it.⁷

⁵ Any theory that does not allow for a way to flesh out the case so that I end up with knowledge of the conditional claim seems too strong and therefore vulnerable to a *reductio ad absurdum* argument that uses this case.

⁶ Proponents of KFK do not have to allow this since it is an independent issue whether beliefs can amount to knowledge when they are based on one pedigree route that meets muster and another that does not.

⁷ On Williamson's approach, warranted assertion is an important heuristic for determining whether someone knows. So there is some reason to think that on Williamson's approach, my case is a successful counterexample to KFK since it is a case in which I am warranted in asserting that if Dana was born on a Tuesday then she was born on a weekday, but I am not warranted in asserting that Dana was born on a Tuesday.

2. Objection 1: Significance and Epistemic Credentials

A defender of KFK might respond by challenging the significance of my case. This might be done by claiming that the key issue concerning KFK is whether a conclusion can be inferentially known on the basis of an inference from one or more beliefs that lack some epistemic credential. The defender of KFK might then assert that for this reason, we should limit ourselves to cases in which all the starting points of a person's inference are claims that the person believes and then focus on whether those beliefs must have various epistemic credentials.

I offer three points in response. First we need to distinguish between two exercises. One involves limiting ourselves to a narrow range of cases. For example, we can do what the defender of KFK recommends, and limit ourselves to cases in which someone believes some claims and goes on to inferentially know something on that basis. We can then ask whether their beliefs in those premises must have various epistemic credentials. A similar sort of exercise involves limiting ourselves to cases in which all of the starting claims are true and someone goes on to inferentially know something on their basis. Here we ask whether the starting claims in these cases must always be claims that the person is justified in believing. The key point though is this: KFK is not limited in this way. Since it is a thesis about the broader genus that subsumes all of these previous cases, testing it involves carrying out a different exercise. This is because KFK says something about all cases in which we start from consideration of some claims (not just the subcases in which we believe those claims, not just the subcases in which those claims are true, etc.), we perform some inferences from those claims, and the result is inferential knowledge of a conclusion. In doing so, it imposes a requirement on the psychological attitude that one must have toward the starting claims of one's inferential reasoning. For this reason, the case I offered is a genuine testcase for KFK.

Second the class of epistemic credentials is not clear, nor is its significance. Consider truth. Often it is not classified as an epistemic credential. After all, false beliefs can be just fine from the epistemic point of view. They can be epistemically justified, epistemically rational, etc. Against this, it might be argued that false beliefs cannot be warranted. But that is contentious.⁸ And even if we grant that false beliefs cannot be warranted, Klein and Warfield do not appeal to the unwarranted nature of the premise-beliefs in their cases. They appeal to the falsity of those beliefs. Importantly this doesn't make their cases any less probative or less interesting. This

⁸ There is a significant literature on this. A good place to start is Michael Huemer, "Logical Properties of Warrant," *Philosophical Studies* 122 (2005): 171-182.

tells against the defender of KFK's view that the only interesting counterexamples to KFK involve premise-beliefs that lack epistemic credentials.

Third, we should ask ourselves whether failing to believe the starting claims in one's inferential reasoning can amount to lacking an epistemic credential that might then explain a failure to know a conclusion that one has deduced. I have suggested that at least when one reasons via conditional proof, it need not. But arguably in other cases it does. Consider someone who competently deduces q in modus ponens fashion, but does so from assuming, rather than believing, p and if p then q . Does this yield inferential knowledge of q ? To keep the focus on belief, suppose the person is warranted in believing p and warranted in believing if p then q (though they don't believe either of these claims) and that each of p and if p then q are true. While things are not entirely clear here, it seems plausible that this person does not end up inferentially knowing q , and that this is so because the person fails to have the right attitude to p and if p then q . If that is correct, then failing to believe one's starting claims might in some cases explain a failure to know an inferred claim. This undercuts the objector's assumption that failing to believe can never itself be a lack of an epistemic credential that might explain a failure to know a competently deduced conclusion.

It is important to notice that the verdict that I just suggested in this last modus ponens case is consistent with the verdict that I am urging in the Dana's Birthday case. A simple picture that delivers both of these verdicts is one on which different combinations of inference patterns and attitudes to one's premises have different results. So the combination of assuming some claims and then reasoning in modus ponens fashion might bar one from knowing a conclusion, while the combination of assuming a claim and then reasoning via conditional proof might not bar one from knowing a conclusion.

3. Objection 2: Smuggled Belief

According to a second objection, I can inferentially know my conclusion only if among my starting points is the belief that all Tuesdays are weekdays. And so my set of starting points includes a belief after all.

There is a simple reply: KFK requires that *all* of the starting points of my inference be claims that I believe. So if I have two starting points, one that I assume and one that I believe, the case can still be a successful counterexample to KFK.

Still it might be proposed that if the belief that all Tuesdays are weekdays is a second starting point, it makes my assumption that Dana was born on a Tuesday superfluous in my coming to know the conditional. It is superfluous, someone might reason, because if I dropped my assumption that Dana was born on a

Peter Murphy

Tuesday, I would still infer from just my belief that all Tuesdays are weekdays to the same conditional claim that if Dana was born on a Tuesday then she was born on a weekday. It now looks like my belief in the conditional is overdetermined. But if that is right, my belief in the conditional is not essentially inferred from my assumption, and so the case is not covered by, and therefore cannot be a successful counterexample to, the KFK principle.

However, there is a crucial difference between the earlier case that I used to illustrate the need for the essentially inferred clause and the Dana's Birthday case. In the earlier case, my belief in the conditional was based on two separate actual sequences: one involved reasoning via a conditional proof, and the other involved someone's testimony. But even if it is true that in the Dana's Birthday case, I must believe that all Tuesdays are weekdays, and it is true that this belief must play a premise role in my conditional proof, there are not two separate actual sequences that lead me to believe the conditional. There is only one. This is so even if – as I am presently conceding – that route must involve both my assumption that Dana was born on a Tuesday and my belief that all Tuesdays are weekdays. For this reason, it is not an overdetermination case. The Dana's Birthday case can therefore be used to test KFK after all.

4. Two Corollaries

My original Dana's Birthday case is just two short steps away from a supercase against KFK. Add two things to it: that it is false that Dana was born on a Tuesday and that I possess no justification or warrant for believing that Dana was born on a Tuesday. Now consider again my assumption that Dana was born on a Tuesday and imagine that I go through the same conditional proof. Now the starting point for my inference does not meet any of the usual conditions on knowledge: I don't believe it is, it is false, and I have no justification or warrant for believing it. Despite all of this, I know the conditional conclusion that I deduce from it.

Here is a second interesting result. Should a foundationalist classify my knowledge of my concluding conditional as basic knowledge or as non-basic knowledge? Basic knowledge is knowledge that is not based on any other knowledge, while non-basic knowledge is based on other knowledge. Is my knowledge of the conditional based on any other knowledge? This takes us back to the smuggling objection. Recall the charge made there was that I must also believe that all Tuesdays are weekdays, and that this belief must function as an additional starting point in my inferential reasoning. While this went unchallenged in my reply to the smuggling objection, there is good reason to think that I need not believe that all Tuesdays are weekdays. That reason comes from the usual lesson

drawn regarding Lewis Carroll's famous regress.⁹ The lesson there is that competent deductive reasoning from some set of claims to a conclusion does not require as a premise in one's reasoning, belief in the linking conditional that features the conjunction of the other premises in one's reasoning in antecedent position and features the conclusion in one's reasoning in consequent position.¹⁰ If this is right, my knowledge of the conditional can proceed from just the assumption that Dana was born on a Tuesday. So it need not be based on any beliefs, and thus on any knowledge, whatsoever. This makes my knowledge of the conditional, basic knowledge. But nonetheless it is inferential knowledge since its source is inference, and not perception, memory, testimony, etc. This makes it an instance of inferential basic knowledge.

I end with a contrast between two big pictures. KFK portrays inference as always in need of serious epistemic help if it is to generate knowledge. This is because KFK claims that all of the starting claims of one's inference must be instances of knowledge, and that the original starting claims back at the origins of one's inference must be known by some other source like perception, memory, testimony, etc. Rejecting KFK results in a very different picture of inference, one that promotes inference to the class of sources that can autonomously generate knowledge without any such help from other sources.

⁹ Lewis Carroll, "What the Tortoise Said to Achilles," *Mind* 4 (1895): 278-280.

¹⁰ This is a bit abbreviated. We need to add that if belief in the linking conditional, if Dana was born on a Tuesday, then Dana was born on a weekday, is not needed then belief in the very closely related claim that all Tuesdays are weekdays is not needed either.