AMERICAN PHILOSOPHICAL QUARTERLY
Volume 57, Number 3, July 2020

KNOWLEDGE FROM KNOWLEDGE

Rodrigo Borges

ABSTRACT

This paper argues that a necessary condition on inferential knowledge is that one knows all the propositions that knowledge depends on. That is, I will argue in support of a principle I call the *Knowledge from Knowledge* principle:

(KFK) S knows that p via inference or reasoning only if S knows all the propositions on which p depends.

KFK meshes well with the natural idea that (at least with respect to deductively valid or inductively strong arguments) the epistemic status of one's belief in the conclusion of an argument is a function of the epistemic status of the attitudes one has to the premises of that argument. One gets what one puts in: S's belief in the conclusion C of an argument A has epistemic status E for S only if S's belief in each of the premises C depends on also has epistemic status E.[1] This is compatible, of course, with the epistemic status of S's conclusion being greater than the epistemic status of the premises on which the conclusion of her reasoning depends. Furthermore, even though KFK applies more straightforwardly to non-suppositional inferences, we can also plausibly apply the principle to sup-positional reasoning if we treat S's assuming for the sake of argument that p,[2] as S assuming for the sake of argument that she knows that p.[3]

Although extremely plausible, philosophers have failed to present a concerted case in support of KFK. Most of their energy goes to articulating and/or defending some form of closure principle for knowledge.[4] What is more, recently versions of this principle have been dismissed as problematic but I think without the reasons in its support been fully appreciated. This paper is an e fort to change the current state of affairs. Its focus is *positive*, however: my goal is to make the best case I can *in support* of KFK.[5] To that end, the paper presents the reasons we have to accept KFK (§1), considers an objection to the principle (§2), and situates the principle in the broader normative context of knowledge norms (§3).

1. THE CASE FOR KFK

I will now present a few reasons we have to accept KFK. For the most part I will not argue in support of the reasons themselves. This is a task for some other occasion. (The exception here is the last reason I will discuss.)

The first reason we have to accept KFK is that it helps explain why, in many circumstances, a reasoner fails to know the conclusion of an otherwise sound argument. We can illustrate this by looking at a famous case from the epistemological literature. Mostly everyone agrees that the protagonist in those cases fails to know at least one of the premises on which her conclusion depends, but discussions of those cases usually focus on what prevents their protagonist from knowing one of the premises. In contrast with this literature, we are here interested

in explaining why the protagonist in those cases fails to know the conclusion of her inference.

Suppose I bought a ticket in a large and fair lottery. Suppose further that my ticket lost but, even though I have not yet checked the results, I nonetheless believe my ticket lost on the basis of my knowledge of the odds. In this situation, were I to reason in the following way, I would not know the conclusion of my argument even though the argument is valid and I have a justified true belief in each of the premises the conclusion depends on:

1. My ticket lost.
2. If my ticket lost, then I will not be able to afford Michael Jackson's mansion.

Therefore,

3. I will not be able to afford Michael Jackson's mansion.

In the above argument, it seems that I do not know I will not be able to a ford Michael Jackson's mansion. If we suppose that I have no other source of money (e.g., inheritance) that would allow me to buy his residence, the natural explanation of why this is the case is the fact that this conclusion depends on something I do not know, namely 1.

Many believe that I have a justified true belief in 1 and 2, but that I nonetheless fail to know those propositions.[6] Since the argument is valid, those people are not free to say that I fail to know 3 because I am not justified in believing 1 and 2, or because they are false. On the other hand, KFK can fill the explanatory gap they face in the following straightforward way: if KFK is true, then the reason why I do not know I will not be able to afford Jackson's mansion is that I do not know all the premises on which this conclusion depends.[7]

This, to be sure, does not mean that we have an explanation of *why I fail to know either 1 or 2*—this is a separate matter. The point that is relevant here is just that, *given* that I fail to know either 1 or 2, and, *given* that KFK is true, one may explain *why I fail to know 3*: because KFK is true, one knows a conclusion of one's reasoning only if one knows all the premises, and I fail to know all the premises in my argument. Moreover, the fact that KFK explains why the latter fact holds and not why I fail to know 1 or 2 is a virtue not a vice, for it does not rule out *a priori* any story about why I fail to know 1 or 2.

What is more, the intuitive appeal of KFK has not passed unnoticed by philosophers from different eras. Aristotle, René Descartes, John Locke, Immanuel Kant, Bertrand Russell, and Timothy Williamson all accept something like KFK. With the exception of Russell and Williamson, the other philosophers accepted something that entails KFK; a principle that required not mere knowledge of the premises on which one's conclusion depends but in fact *a priori* knowledge of those premises.[8] So, KFK (if true) not only explains why otherwise sound inferences fail to yield knowledge of their conclusion but also plays a prominent role in the distant (and not-so-distant) history of epistemology. Of course, tradition is not always right, but the wide acceptance of KFK should, I think, carry some weight when we consider the principle.

An increasingly popular account of evidence also supports KFK. Williamson has argued that one's evidence is constituted solely by what one knows. According to this view *E*vidence is identical to *K*nowledge (i.e., E=K).[9] Since evidence is what justifies inferential knowledge, according to E=K, knowledge is what justifies knowledge.[10] Thus, only reasoning from what is known justifies knowledge. KFK is vindicated because E=K entails that only known relevant premises justify knowledge. KFK does not entail E=K, however, for KFK merely requires that evidence relevant to one's conclusion be known, while E=K requires that *all* evidence be known — whether relevant or irrelevant to one's conclusion. (Someone who rejected E=K could hold the view

that, although *knowledge* of a conclusion requires knowledge of all the premises in one's reasoning (i.e., inferential knowledge conforms to KFK), *other* epistemic states (e.g., justified belief in one's conclusion) can be acquired via reasoning even if one does not know all the premises involved). As a result, KFK is *at least* as plausible as E=K. Of course, this will not move anyone who does not like E=K.[11]

Another reason to accept KFK is that it is compatible with (if not entailed by) different analyses of the concept *knowledge*. The point here is not to endorse those analyses, but, rather, to show that KFK gains plausibility from being compatible with different analyses of this concept. Furthermore, I do not mean to suggest that *knowledge can* or even *should* be analyzed. Maybe the task of analyzing this concept *is* a waste of time.[12] Still, many think this is not the case and view the task of analyzing *knowledge* worth their while. Some of those analyzes are committed to KFK.

For example, Peter Klein's account of knowledge as undefeated justified true belief arguably entails KFK. According to Klein's view, one knows that p if and only if one has a justified true belief that p and there is no truth d such that the conjunction of the propositions in one's evidence, e, and d, fails to justify one in believing that p. If there is such a truth, then one's justification for believing is *defeated*, and one fails to know that p.[13]

KFK is entailed by this analysis of knowledge because, if one fails to know a premise, x, on which one's conclusion depends, then there is a truth ¬k (namely, "S does not know that x") that is such that the conjunction of ¬k and those premises fails to justify one's belief in the argument's conclusion.[14] More specifically, ¬k defeats one's justification because either one does not believe that x, x is false, or x is not justified for one; the truth of any of these claims entails ¬k and is sufficient, according to Klein, to defeat one's justification.

Additionally, Ernest Sosa's virtue epistemology can easily incorporate KFK. According to Sosa, knowledge is apt belief.[15] Roughly, a belief is apt if and only if it is not only true, but true *in virtue of* it being the output of a virtuous cognitive process. Within this framework we can say that the belief in a conclusion is aptly formed (i.e., it is a case of inferential knowledge) only if the agent aptly believes all the premises on which her conclusion depends. This is a welcome addition to Sosa's epistemology, given the way he sometimes treats the original Gettier cases.[16] According to Sosa, Smith in the original cases has a justified true belief that falls short of knowledge *because the conclusion of Smith's reasoning depends essentially on a false belief*. Although this diagnosis of what is going wrong in the original Gettier cases is (as we will see below) a popular one, its adequacy is limited: although the requirement that one's conclusion depend only on *true* premises is clearly necessary, it is just as clearly not sufficient, for reliance on a true but *un*justified premise cannot bestow justifiedness on one's conclusion. But, given the guiding intuition we started with (i.e., the idea that the epistemic status of one's belief in a conclusion is to a great extent determined by the epistemic status of one's belief in each of the premises on which one's conclusion depends), it follows that reliance on aptly formed beliefs is necessary for apt belief in one's conclusion. So, although not explicitly stated by Sosa, a natural improvement of his diagnosis of the original Gettier cases leads straight to KFK.[17]

As it might already be obvious to the reader, if KFK is true then we also have a straightforward explanation of the Gettier phenomenon (i.e., the fact that get-tiered agents fail to know although they have a justified true belief). Because KFK is a general principle about inferential knowledge for which we have in-dependent reasons to believe, the principle makes for an even more plausible explanation of this phenomenon.[18]

To illustrate this point, let us apply KFK to Gettier's Case I. In that case Smith reasons thus:

1. Jones has ten coins in his pocket and Jones will get the job.

   Therefore,

2. The man who will get the job has ten coins in his pocket.

Although Smith is justified in believing 1, he does not know it since it is false. If KFK is true, then Smith fails to know 2 because he fails to know 1, and 2 depends on 1.

One might think that appealing to KFK cannot be the whole story about what goes awry in Gettier cases because one thinks that there are non-inferential Gettier cases. Allegedly, Roderick Chisholm's sheep case and Alvin Goldman's barn-façade case are examples of non-inferential Gettier cases.[19] If this is right, then KFK would not be the whole story about the Gettier phenomenon for it applies only to cases of explicit reasoning or inference. This objection is problematic, however. Some have persuasively argued that Chisholm's case is not re-ally non-inferential and that its protagonist derives a justified true belief from a justified but false belief;[20] and others have argued that Goldman's case is in fact a case of knowledge.[21] Given those doubts, one might argue that we should (at least temporarily) suspend judgment as to whether those cases are in fact Gettier cases since there is no consensus around this issue. However, I am now inclined to think that this reaction is too strong.[22] Even if those cases are in fact Gettier cases, it is not the case that KFK has to explain *all* Gettier cases in order to be attractive—there may be more than one failure in different Gettier cases, as some have argued (e.g., Pritchard 2005)—, it is more than enough that KFK explains what goes wrong in the original cases and in all the cases that share their structure.[23]

This brief overview shows that KFK enjoys strong intuitive and theoretical support. Given such a strong case in its favor, this principle should play a more prominent role in epistemological discussions than it currently does.

2. Knowledge from Non-Knowledge?

The result of the previous section is that KFK enjoys great support. I take this to show that our commitment (explicit or not) to KFK is central to our epistemic commitments more generally. As a result, only a reason of considerable weight should be allowed to undermine our commitment to this principle. In this section I will look at a recently proposed counterexample to KFK. I will argue that, once everything is said and done, the counterexample fails. The result will be that our commitment to KFK is well justified.[24]

Alleged cases of knowledge from non-knowledge[25] seem to show that knowledge of the premises one's conclusion depends on is *not* necessary for knowledge of said conclusion. In those cases, the protagonist allegedly comes to know some truth even though her conclusion depends on a false premise.[26]

Here is one such case.[27] Ted printed 100 copies of the handout for his talk and, after carefully counting 53 people in attendance, he concludes that 100 hand-outs are enough. However, Ted double-counted one person who moved during his counting of heads, making his premise that there are 53 people in attendance false.

This, and other examples like it, are supposed to show that sometimes one may acquire inferential knowledge even though one does not know all the propositions on which one's conclusion depends. There are at least two different strategies we can use to deal with those cases. Those strategies differ on how much they concede to the KFK denier.

One might insist that there is knowledge in cases such as Warfield's but argue that this knowledge does not depend on the false proposition. This reply relies on a precise account of epistemic dependence. On the

other hand, one might argue that the protagonist in those cases does *not* know the conclusion of her inference. This view says that Ted's conclusion depends on a premise he does not know but denies that Ted knows his conclusion. Rather, his belief has some other epistemically good-making property (e.g., it is probably true, given his evidence). Although I personally favor the latter view, I sometimes vacillate between assigning knowledge and assigning ignorance to the protagonist of cases such as Warfield's. That is why I offer both replies—just so that I can have (some) comfort every day of the week. I will discuss each reply in turn.

Arguing that Ted's knowledge of the conclusion does not depend on his false belief requires showing that

(t) There are enough handouts

does not depend on

(f) There are 53 people in attendance.

I have so far relied on an intuitive understanding of the notion of dependence. It is now time to be more precise about this notion. We will see that, under a plausible interpretation of what it takes for x to depend on y, this relation does not hold between t and f.

Epistemically speaking, S's belief that p depends on S's belief that q in a case C if and only if (i) $P(p \mid q) > P(p)$ (i.e., p depends on q only if the probability of p given q is higher than p's unconditional probability),[28] (ii) the unconditional probability of p, P (p), is not sufficient for S to be justified in believing p, and (iii) S's belief in q partially explains why she believes that p.

Condition (i) seems necessary because the conclusion of one's argument de-pends on some premise only if this premise "counts in favor" of said conclusion, or (to put it another way) the conclusion depends on some premise only if the conclusion, conditional on the premises, is more likely to be true than if considered by itself. (ii) is necessary because q might count in favor of p in this sense but still no count in favor of p enough to justify one's belief in p. Condition (iii) seems to be necessary as well because we usually think that a belief depends on another only if the first exists in part because or in virtue of the later also existing.[29] Thus, given an interpretation of dependence based on (i)-(iii), x depends on y only if y justifies and partially explains x. Here is a more natural way to capture this idea:

> (Dependence) S's belief that p depends on her belief that q in C iff (i) without q, S's evidence set does not justify her belief that p in C and (ii) the fact that S believes that q in C at least partially explains why she believes that p in C.[30]

Now, apply Dependence to Warfield's case. Although the fact that Ted believes that there are enough handouts is in part *explained* by the fact that he believes that there are 53 people in attendance, it is not the case that *without* this belief his evidence set *does not* justify him in believing that he has enough handouts. That is to say, when we substitute "There are 53 people in attendance" and "There are enough handouts" for q and p, respectively, in Dependence, we see that, although condition (ii) is satisfied, condition (i) is not. The upshot is that Ted's knowing that there are enough handouts does not depend on his false belief that there are 53 people in attendance. If this result can be coupled with an account of how Ted can know his conclusion without depending on his false belief for this knowledge, then we will have mounted a strong enough case in defense of KFK against the threat coming from alleged cases of knowledge from non-knowledge. Here is such an account of how Ted knows that there are enough handouts with-out dependence on his false belief.

The first thing to note is that Ted has, roughly, twice as many handouts as there are people in attendance. Since, as the case is described, there are no obstacles between

Ted and the attendees and since his vision is working properly, it can be plausibly argued that Ted's evidence justifies him in believing he has enough handouts *even before* he finishes counting the number of people in the room. After all, as one counts a finite number of objects, all of which are in one's unobstructed field of vision, one not only acquires increasingly more precise in-formation about how many objects there are but one also acquires increasingly more precise information about how many objects remain unaccounted for. So, if I want to take twenty bottles of my favorite beer to the party and there are forty bottles of that beer in my fridge, I know I have enough beer bottles to take to party even before I count them all, if I have counted, say, ten bottles. After having counted ten beer bottles (and thereby coming to know there are ten bottles of my favorite beer in the fridge), I can also come to know, by abduction from how things inside the fridge look to me, that there are more beer bottles I haven't yet counted than the number of bottles I need (i.e., ten). This abductive step is justified by my background knowledge (e.g., I know that the bottles of my favorite beer look a certain way and that the bottles in my fridge look that way).[31] I am therefore justified in believing I have enough bottles of my favorite beer even though I have not counted all of them. So much so, that it seems to me that no one could plausibly fault me for having stopped counting the bottles when I did. In fact, in circumstances as mine, counting all beer bottles would be a waste of my time.

Similarly, after Ted counts, say, forty people, he comes to know, by abduction from how things look inside the room, *that there are less than sixty people left*. This abductive step is justified for Ted because of his background knowledge: he knows what people look like and where they tend to be located in situations such as his, and so on.[32] Thus, Ted is justified in believing that his one hundred handouts are enough. In fact, his knowledge that there are forty people and that there are less than sixty people left *entails* that he has enough handouts, if we allow "there are enough handouts," in this context, to mean "the number of handouts is greater than or equal to the number of people in attendance." Although Ted did not stop counting until (after) he attributed a number to every-one in attendance, I do not think it is implausible to say that we could not have plausibly faulted him if he had stopped counting at forty and concluded that he had enough handouts. In fact, I believe that what makes this hypothetical judgment plausible is precisely the fact that Ted knows he has enough handouts even before he finishes counting everyone in the room.

Someone might complain that the claim that Ted knows that there are less than sixty people left in the room after he counted forty is implausible on the grounds that knowledge entails belief and Ted did not explicitly formed the belief that there are less than sixty people left in the room. This objection does not work, for many (if not most) of our beliefs are *not* consciously formed. In fact, it is widely known that humans non-explicitly count (and thereby acquire knowledge about) small numbers of objects—this is the sudden and subconscious mental process known as "subitizing."[33]

It is important to note that the number of people Ted needs to count before the abductive step becomes justified for him might well be vague. It seems clear to me that, after having counted one single person, the claim that there are less than ninety-nine people left is not justified for Ted, and neither is the claim that he has enough handouts. It seems just as clear to me that, after having counted fifty-two people, the claim that there are less than forty-nine people left *is* justified for Ted (remember that there's only 52 people in attendance). What is less clear is the answer to the question "When exactly a claim of this

type goes from not being justified to being justified for Ted?" On the other hand, I think that, after having counted forty people, it is fairly clear that the claim that there are less than sixty people left is justified for Ted. Regardless of where one chooses to draw a line here, the point remains that, by the time Ted actually draws the conclusion that he has enough handouts, he is justified in believing that he has enough handouts independently of his false belief that there are fifty-three people in attendance.[34]

This diagnosis of Ted's case can be generalized into a diagnosis of all the other cases of knowledge from non-knowledge in the literature in the following way. Given Dependence, and given the general structure of knowledge from non-knowledge cases (i.e., a known conclusion that seems to be based essentially on a premise that is not known), if C is a case of knowledge from non-knowledge, then S's known conclusion in C *does not depend* on the relevant unknown premise, in the sense of "depends on" described by Dependence;[35] instead, S's knowledge of the conclusion depends on something else she knows. Even though the "something else she knows" I am referring to here will change from case to case (in part because agents in different cases have different evidence sets), the general diagnosis of those cases remains. I will not apply this diagnosis to *all* cases of knowledge from non-knowledge in the literature. But I will discuss one other case as a way of illustrating further how we should apply this diagnosis. Liz suspects her son has a fever and decides to measure his temperature using her reliable mercury thermometer. She competently uses the thermometer to measure her son's temperature and comes to believe it reads 40 degrees Celsius. Since Liz knows that a temperature above 37 degrees Celsius indicates fever, she concludes that her son has a fever. Although Liz is right that her son has a fever, she misread the thermometer—it accurately showed 39.8 degrees Celsius, not 40 degrees Celsius.[36] Because Liz knows that her son has a fever even though she inferred this from her false belief that the thermometer reads 40 degrees Celsius, this is a knowledge from non-knowledge case.

Now, according to the general diagnosis of those cases I offered above, Liz's knowledge that her son has a fever does not depend on her false belief that the thermometer reads 40 degrees Celsius. True, the fact that she has this false belief helps *explain* why she also believes her son has a fever; yet, Liz's false belief is not what *justifies* her in believing her son has a fever. What justifies her in believing her son has a fever is her knowledge *that the thermometer shows a temperature above 37 degrees Celsius*, something she learns by looking at the thermometer. Furthermore, this is no idiosyncrasy on Liz's part, for mercury thermometers used to measure body temperature all have a very salient red line marking the threshold of fever (i.e., 37 degrees Celsius), which makes knowing that a thermometer shows a temperature above 37 degrees pretty common in a case of fever. The upshot is that Liz knows her son has a fever, but what justifies her in believing that he does is not her false belief that the thermometer reads 40 degrees Celsius.

I conclude, then, that there is much that can be said in favor of KFK and against the claim that particular cases show that knowledge from non-knowledge is possible.

3. The Knowledge Norm of Theoretical Inference

Although I believe something like Dependence is correct, I sometimes do feel like saying that Ted's conclusion depends (in a sense of "depends" different from Dependence) on his false belief, and that the intuition that Ted knows that 100 handouts are enough is *not* probative of him actually knowing that this is the case. In this section I try to make sense of this way of looking at cases of knowledge from non-knowledge. I conclude that

accounting for those cases in this way makes at least as much sense as accounting for them in the way I suggested above. The main hurdle we need to clear in order to make this story plausible is showing that it is not too crazy to think that the protagonist of alleged cases of knowledge from non-knowledge doesn't in fact know the conclusion of her argument. I will do that by arguing that this conclusion makes perfect sense when we embed it in a framework that is independently motivated—the normative framework of the knowledge first account of knowledge.

The general idea is simple enough. Take Ted's case again. Ted does not satisfy KFK and therefore does not know that 100 handouts are enough. We then couple this assessment of the case with an error theory which explains away the intuition that Ted knows: although Ted does not know that 100 handouts are enough, Ted is justified in believing they are, for he has an *excuse* for believing the way he does. After all, he did count the number of people carefully and he did not miss the right number by much. But, having a justified true belief is not the same as knowing. So, even though Ted's belief is well justified and, from his point of view, "as good as knowledge," he does not know that 100 handouts are enough. I will expand below on this theme and show how it fits well with a variety of knowledge norms. In the meantime, let me say this about all alleged cases of knowledge from non-knowledge.[37]

According to the view being articulated here, in all alleged cases of knowledge from non-knowledge, the protagonist fails to know the conclusion of her reasoning, but it is highly probable, on the protagonist's evidence, that she knows the false premise of her argument. And, since "I know that p" is highly probable on the protagonist's evidence, it is *reasonable* for her to believe that p and she has an *excuse* to believe that p.[38] However, it being excusable (because reasonable) for one to believe each of one's premises is necessary, but obviously not sufficient, to generate knowledge of one's conclusion. Thus, the protagonists of cases of alleged knowledge from non-knowledge have inferentially *reasonable* (or justified) belief, rather than knowledge of the conclusion of their inference. As we will see below, this account of what it is to have an excuse can be modeled after the knowledge norms of practical inference and assertion.[39]

So, even though it is false that the protagonists of knowledge from non-knowledge cases knows the conclusion of their reasoning, we can accommodate the related intuition *that it is reasonable for them to believe as they do*. Moreover, we are liable to mistake this fact for the non-fact that the protagonist in a case of alleged knowledge from non-knowledge knows, because (i) it is reasonable for her to believe she knows all beliefs on which her conclusion depends; (ii) there is virtually no difference, from the protagonist's point of view, between knowing that p and believing that p is true; and, (iii) when we assess philosophical cases we tend to put ourselves into the protagonist's shoes and let this influence our assessment of the case itself.

One might object that this approach to alleged cases of knowledge from non-knowledge blurs the distinction between those cases and Gettier cases. In both Gettier cases and knowledge from non-knowledge cases the protagonist arrives at a justified true belief that is not knowledge in virtue of a deduction from a false but justified belief. However, these are necessary conditions on a case being a Gettier case, not necessary *and* sufficient conditions for a case to Gettier case.[40]

An obvious candidate for such a further condition is that a case is a Gettier case only if it is fairly obvious that the protagonist in that case fails to know. If that is correct, then alleged cases of knowledge from non-knowledge cannot be Gettier cases, for, as I

suggested above, it is not "obvious" that there is no knowledge in those cases.

In this way of thinking about things, the normativity of KFK stems from the following *Knowledge Norm of Theoretical Inference*:

(KNI) One must: infer only from what one knows.

KNI applies most directly to non-suppositional reasoning, but it could, I think, be extended to suppositional reasoning like this: for S to suppose, for the sake of argument, that p, is for S to suppose, for the sake of argument, that she knows that p.

There are a few structural similarities between KNI and the *knowledge norm of practical inference*.[41]

(KNC) One must: act only on what one knows.

If one thinks that one should not act unless one knows all the propositions one uses as reasons for one's actions, then one thinks that rational action requires good reasons and that reasons are good only if they are known. This is, *mutatis mutandis,* the idea behind KNI: rational belief requires good reasons and reasons are good only if they are known.[42]

First, let us apply KNI to the lottery case. We saw that it makes perfect sense for someone to say that, since I do not know that my ticket lost, I should not infer that I will not be able to a ford Michael Jackson's mansion. In the practical case, since I do not know that my ticket lost, it is appropriate for someone to criticize me for throwing the ticket in the garbage bin. If the knowledge norm of action is true, then my throwing the ticket in the bin is inappropriate, for I base that action on something I don't know, thus breaking the knowledge norm of action. Similarly, if KNI is true, we can explain why my *belief* in the conclusion that I will not be able to a ford Michael Jackson's mansion is inappropriate for, by so believing, I break that norm.

The same applies to the original Gettier cases and to the knowledge from non-knowledge cases. Because Smith and Ted do not know the premise of their inferences, they should not have *believed* their conclusion on its basis. This is what KNI says. Interestingly enough, if we suppose that Smith leaves the office because he believes falsely that Jones will get the job, then KNC yields the result that Smith's leaving is not justified. The same is true of Ted in the handout case. He breaks KNI for believing that 100 handouts are enough on the basis of his false belief that there 53 people in attendance; but he also breaks KNC when he acts on the basis of that false belief and distributes his handout.

No doubt, some will complain that this is a harsh assessment of the protagonist in those cases. "Sure," they will say, "there is a clear sense in which believing and acting on the basis of what they fail to know *is* appropriate; it is a problem for both KNI and KNC that they can't capture this propriety." I think this criticism raises an important point and addressing it will discharge the promise I made in the previous section that there is something epistemically good about Ted's belief that 100 handouts are enough, even though he does not know that that is the case. However, in order to do that, we will have to take a closer look at Williamson's account of excusable breaches of the *knowledge norm of assertion*:[43]

(KNA) One must: assert only what one knows.

KNA determines proper assertion: S asserts that p *properly* if and only if S knows that p.[44] But, sometimes one asserts things one does not know and those assertions seem, in some sense, *proper*. Following Williamson (2000, p.257), we can account for those cases if we distinguish between *S having a warrant to assert that p* and it *being reasonable for S to assert that p*.[45] He considers the following case: it is winter, and it looks exactly as it would if there were snow outside, but in fact that white stuff is not snow but foam put there by a lm crew of whose existence I have no

idea. I do not know that there is snow outside because there is no snow outside, but it is quite reasonable for me to believe not just that there is snow outside but also that I know that there is; for me, it is to all appearances a banal case of perceptual knowledge. Surely it is then reasonable for me to assert that there is snow outside.

The case is consistent with KNA because, if I am entitled to assume KNA and if it is reasonable for me to believe that I have a warrant to assert that there is snow outside, then it is reasonable for me to assert that there is snow outside. The consequence is that we can use KNA to explain the *reasonableness* of my assertion in the snow case, while, at the same time, denying that this assertion is *warranted* for me.

The idea here is that knowledge gives you warrant to assert while reasonable-ness makes your assertion *blameless*. One's assertion that p is reasonable and one asserts that p blamelessly if "I know that p" is probable enough on one's evidence.

We are now in a position to reply to the objection that it was implausible for us to criticize Smith and Ted for breaking KNI and KNA, because there was a clear sense in which their belief and action were appropriate.

It is reasonable, in the sense just introduced, for both Smith and Ted to believe and act the way they did. Although they break KNI and KNA, they are blameless for doing so because it is reasonable for them to believe that they have a warrant to believe and act the way they do. In other words, Smith and Ted are at fault for failing to know the relevant propositions and thereby breaking KNI and KNA, but they can be excused for breaching those norms in virtue of the fact that it is probable, given their evidence, that they know those propositions.

One familiar complaint about knowledge norms is that, even if we are not skeptics, they are too demanding. Some say that we are not always in a position to know whether we know the premises in our reasoning and that we should settle for a weaker epistemic standard such as justified belief, for we are relatively good at finding out whether or not we are justified in believing the premises of our inferences.

The usual reply to this kind of objection to knowledge norms can also be used in defense of KNI. First, KNI does not require one to know that one knows the premises to one's reasoning—first order knowledge (i.e., *knowing* the premises) is enough. Second, the alleged fact that KNI is broken frequently does not show KNI is not the epistemic rule governing epistemically proper inference; the fact that many people break the rule "one must: stop at stop signs" does not show that this rule is not what governs proper driving behavior at intersections with stop signs.

KNI would be too demanding only if we knew too little, for, if that were the case, then we would have very little to work with in our inferences. However, since we clearly sometimes reason properly, we might plausibly take that to confirm KNI. Skeptics about knowledge will, of course, demur, but we have as much reason to think the skeptic is right about this as we have to think that there is no knowledge at all. While skepticism is an important philosophical question, it presents no special challenge to KNI.

What is more, epistemic norms that are weaker than KNI cannot account for the propriety of this criticism in a way that is as straightforward as KNI. Although the protagonist of the lottery case fails to know one of his premises, his belief is true and justified. But, if the normative standard of proper inference was given by the *justified true belief norm*,

> (JTBNI) One must: infer only from what one truly and justifiably believes,

then we should expect my criticism of the protagonist in the lottery case to "sound

bad," for he has a justified true belief in the premise of his reasoning. A *true belief norm* fares even worse.

(TBNI) One must: infer only from what one believes truly

If such a norm were correct, it would not only predict that our criticism of Smith was mistaken, but it would also entail that one could acquire inferential knowledge even if the sole reason one believed the premises was because of a coin toss. Hence, KNI has an advantage over weaker epistemic norms.

What is more, it seems that someone who accepts the *knowledge norm of belief*,

(KNB) One must: believe only what one knows,

is committed to KNI: if one's belief that p depends essentially on one's belief that q and one fails to know q, then one has not only violated KNB but also KNI.

In sum, to the extent that we care about knowledge and to the extent that we think that knowledge should inform how we assess the correctness of inferences, KNI can do a lot to accommodate those values.

University of Florida

NOTES

This paper benefited from the generosity and support of Peter Klein. Claudio de Almeida, Duncan Pritchard, Ernie Sosa, Timothy Williamson, Mike Veber, and Stephen Hetherington discussed different drafts with me. Many thanks to all of them. Thanks also to two reviewers for APQ for their careful and helpful work. The research for this paper was partially funded by FAPESP.

1. *Nota bene*: KFK *does not* require *all* knowledge to be inferential—it requires all *inferential* knowledge to be based on other knowledge. Thus, KFK is (in principle) compatible with foundationalism, coherentism, and infinitism. Thanks to a reviewer for APQ for prompting me to clarify this issue.

2. Since nothing here hangs on it, I use "argument," "inference," and "reasoning" interchangeably.

3. This addresses a worry raised by Murphy (2013).

4. As in Dretske (2005) and Hawthorne (2005).

5. This paper uniquely explores different reasons in support of KFK; e.g., section 3 argues that KFK fits the broader normative framework of knowledge norms. I discuss competing views on KFK in endnotes xxv, xxvii, xxxv, and in Borges (Forthcoming).

6. Among others, Lewis (1996), Williamson (2000), Hawthorne (2004), Pritchard (2008), and McGlynn (2013). But, see Nelkin (2000), and Bird (2007) for a dissenting opinion.

7. Lottery cases falsify the *no-false-grounds* account of inferential knowledge. This view, articulated in Clark (1963), Harman (1980), and Lycan (2006) seems to assume that justified belief in *true* premises would be enough to yield knowledge of the conclusion. Lottery cases, on the other hand, strongly suggests that I *do* have a justified true belief in all the premises, but no knowledge of the conclusion.

8. Compare Aristotle (1994, 72a25–30, p.101–2), Descartes (1985, p.15), Locke (1975, BK IV, ch. II, sec. 7), Kant (1950, p.301), Russell (1912, p.132–3), and Williamson (2007, p.145).

9. Williamson (2000, chapter 9).

10. Although Williamson (2000, ch.9) intends E=K to be understood in a way that entails KFK (i.e., given E=K, an inferential belief is justified only if it is supported by evidence/knowledge), some (e.g., Bird (2007)) have suggested that one should understand E=K in a way that does not necessarily entail

KFK (i.e., every known proposition is capable of justifying some inferential belief, but not all inferential belief is justified by evidence). Thanks to a reviewer for APQ for discussion here.

11. See Greenough and Pritchard (2009) for discussion.

12. As in Williamson (2000).

13. See Klein (1981, p.148–166).

14. See Klein (2008, 2012), Borges (2016).

15. See Sosa (2007, 2011).

16. Compare Sosa (2007, p. 95–107).

17. This generalizes to *no-false-grounds* views, and it shows how KFK improves on them. See endnote viii.

18. Compare Armstrong (1973, p.152).

19. See Chisholm (1977) and Goldman (1976).

20. For example, Lycan (2006).

21. Sosa (2007) and Gendler and Hawthorne (2005).

22. I argued for it in Borges (2017).

23. Shope (1983, p. 4).

24. My take on Warfield's case, although compatible with that of Ball and Bloome-Tillmann (2014), Montminy (2014) and Schnee (2015), is also importantly different in that it does not rely merely on our intuitive judgments about particular cases. Below I support KFK with an independently motivated argument about epistemic dependence.

25. Warfield (2005); Klein (2008); Fitelson (2010).

26. Coffman (2008) and Luzzi (2010) present cases in which one allegedly acquires knowledge from a gettiered premise. Although I do not discuss their cases here, my treatment of knowledge from falsehood cases extends straightforwardly to those cases, since gettiered beliefs are not items of knowledge.

27. From Warfield (2005).

28. As in Williamson (2000, ch.10), here "P( )" refers to the *evidential probability* of *contingent* truths.

29. See Sosa (2007, chapter 5).

30. Cases of evidential overdertermination instantiate Dependence in an interesting way. If S can come to know that p via inference from either q1 or q2, then, according to Dependence, p does not depend on q1 or q2. This fits our intuitive grasp of the idea of something depending on something else: if I can get $1 to buy a lollipop from either Mary or Bob, then my buying of the lollipop does not depend on either one of them giving me $1. Of course, S knows that p and I get a lollipop only if at least one of the overdetermining paths obtains. I discuss Dependence further in Borges (Forthcoming).

31. This step is clearly *safe*, since I could not have easily been wrong about the number of bottles left in the fridge.

32. This step, too, is *safe*,—Ted could not easily be mistaken about the number of people left.

33. See Kahneman (2011) for a dual process account of bias and heuristics in general, including subtizing.

34. This reply to Ted's case is importantly different from the view discussed and rejected by Luzzi (2014, p. 264–5). Firstly, Dependence distinguishes *epistemic* and *explanatory/causal* dependence us-

ing conditions (i) and (ii), respectively—something Luzzi says is "absent" from accounts of that case. Secondly, Luzzi takes Ted's relevant evidential basis to be the proposition that there are *approximately* fifty-three people in attendance. This, Luzzi claims, prevents responses to Ted's case from generalizing adequately. Since I identify a different sub-set of Ted's evidence that is sufficient to justify his conclusion, my approach to this case does not suffer from the difficulty identified by Luzzi.

35. That is, the unknown premise helps explain why S believes the conclusion of her argument, but it fails to epistemize this conclusion.

36. Modified from Hilpinen (1988).

37. Some proponents of knowledge norms have recently argued that cases like Handout falsify KFK (e.g., Hawthorne (2017) and Littlejohn (2016)). They think that false beliefs may cause inferential knowledge provided, say, that those beliefs are safe (and that they meet some further conditions). But to the extent that those philosophers take Ted's inferential knowledge to depend *causally*—if not *evidentially*—on his false belief, they agree with the *spirit* (but not the letter) of Dependence. What is more, to the extent that they accept the framework of knowledge norms, I believe they should also accept the knowledge norm of theoretical reasoning I present below. Thanks to a reviewer for APQ for prompting me to address this issue.

38. See Williamson (2011) and Williamson (2015).

39. See Hawthorne and Stanley (2008) and Williamson (2011).

40. See Shope (1983, p.4).

41. Compare Hawthorne and Stanley (2008, p.577).

42. See Unger (1975, p.197–214), and Williamson (2000, chapter 9).

43. Williamson (2000, ch.11)

44. See also Moore (1993) and Unger (1975) for earlier statements of KNA.

45. DeRose (2002, p. 180) and DeRose (2009, p. 94–5) make a similar distinction—endorsed by Weiner (2006) and Benton (2011). Lackey (2007) objects to the distinction.

REFERENCES

Aristotle, and J. Barnes ed. 1994. *Posterior analytics,* 2nd ed. (Oxford, UK: Oxford University Press).

Armstrong, David Malet. 1973. *Belief, truth and knowledge* (Cambridge, UK: Cambridge University Press).

Ball, Brian, and Michael Bloome-Tillmann, M. 2014. "Counter closure and knowledge despite falsehood," *Philosophical Quarterly*, vol. 64, pp. 552–568.

Benton, Matthew. 2011. "Two More for the Knowledge Account of Assertion," *Analysis*, vol 71, no. 4 684–687.

Bird, Alexander. 2007. "Justified Judging," *Philosophy and Phenomenological Research*, vol. 74, no. 1, pp. 81–110.

Borges, Rodrigo. 2016. "A Failed Twist to an Old Problem: A Reply to John N. Williams," *Logos and Episteme*, vol. 7, no. 1, pp. 75–81.

———. 2017. "Inferential Knowledge and the Gettier Conjecture," in *Explaining Knowledge: New Essays on the Gettier Problem*, ed. Rodrigo Borges, Claudio de Almeida, and Peter D. Klein, pp. 273–291. (Oxford, UK: Oxford University Press).

———. Forthcoming. *Knowledge Dependence.*

Chisholm, Roderick. 1977. *Theory of knowledge* (Upper Saddle River, NJ: Prentice Hall).

Clark, M. (1963). "Knowledge and Grounds: A Comment on Mr. Gettier's Paper," *Analysis*, vol. 24, no. 2, pp. 46–48.

Coffman, E.J. 2008. "Warrant Without Truth?" *Synthese*, vol. 162, no. 2, pp. 173–194.

DeRose, Keith. 2002. "Assertion, Knowledge, and Context" *The Philosophical Review*, vol. 111, pp. 167–203.

DeRose, Keith. 2009. *The Case for Contextualism* (Vol. 1). (Oxford, UK: Oxford University Press).

Descartes, Rene. 1985. "Rules for the Direction of the Mind," in *The Philosophical Writings of Descartes*, ed. Dugald Murdoch, John Cottingham, and Robert Stoothoff, vol. I. (Cambridge, UK: Cambridge University Press).

Dretske, Fred. 2005. "Is Knowledge Closed Under Known Entailment? The Case Against Closure," in *Contemporary Debates In Epistemology*, ed. Mathias Steup, and Ernest Sosa. (Malden, MA: Blackwell).

Fitelson, Branden. 2010. "Strengthening the Case for Knowledge from Falsehood." *Analysis*, vol. *70* no. 4, pp. 666–669.

Gendler, Tamar Szabo, and John Hawthorne. 2005. The Real Guide to Fake Barns. *Philosophical Studies*, vol. 124, pp. 331–352.

Goldman, Alvin. 1976. Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, vol. 73, no. 20, 771–791.

Greenough, Patrick, and Pritchard, Duncan. (Eds.). 2009. *Williamson On Knowledge*. (Oxford, UK: Oxford University Press).

Harman, Gilbert. 1980. Reasoning and Evidence One Does Not Possess. *Midwest Studies in Philosophy*, vol. 5, pp. 163–182.

Hawthorne, John. 2005. The Case for Closure, in *Contemporary Debates in Epistemology* ed. Mathias Steup, and Ernest Sosa (Malden, MA: Blackwell).

Hawthorne, John, and Jason Stanley. 2008. Knowledge and Action. *Journal of Philosophy*, vol. 105, no. 10, 571–590.

Hilpinen, Risto. 1988. Knowledge and Conditionals. *Philosophical Perspectives*, vol. 2, pp. 157–182.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. (New York, NY: Farrar, Straus, and Giroux).

Kant, Immanuel. 1950. *Critique of Pure Reason* (N. K. Smith, Ed.). (New York, NY: Macmillan).

Klein, Peter D. 1981. *Certainty: A Refutation of Scepticism*. (Minneapolis, MN: University of Minnesota Press).

Klein, Peter D. 2008. "Useful False Beliefs," in *New Essays in Epistemology*, ed. Quentin Smith, pp. 25–63 (Oxford, UK: Oxford University Press).

Klein, Peter D. 2012. "What Makes Knowledge the Most Highly Prized Type of Belief?" in *The Sensitivity Principle in Epistemology*, ed. Tim Black, and Kelly Becker, , pp. 152–169 (Cambridge, UK: Cambridge University Press).

Lackey, Jennifer. 2007. "Norms of Assertion," *Noûs*, vol. 41, no. 4, pp. 597–626.

Locke, John. 1975. *An Essay Concerning Human Understanding* (P. H. Nidditch, Ed.). (Oxford, UK: Oxford University Press).

Luzzi, Federico. (2010). "Counter-closure," *Australasian Journal of Philosophy*, vol. 88, no. 4, pp. 673–683.

———. (2014). "What Does Knowledge-Yielding Deduction Require of Its Premises?" *Episteme*, vol. 11, no. 3, 1–15.

Lycan, William. 2006. "The Gettier Problem," in *Epistemology Futures*, ed. Stephen Hetherington, pp. 148–168. (Oxford, UK: Oxford University Press).

Montminy, Martin. 2014. "Knowledge Despite Falsehood," *Pacific Philosophical Quarterly*, vol. 44, no. 3–4, pp. 463–475.

Moore, G. E. 1993. "Moore's Paradox," in *G. E. Moore: Selected writings*, ed. Thomas Baldwin (New York, NY: Routledge).

Murphy, Peter. 2013. "Another Blow to Knowledge from Knowledge," *Logos and Episteme*, vol IV, no. 3, pp. 311–317.
Pritchard, Duncan. 2005. *Epistemic luck.* (Oxford, UK: Oxford University Press).
Russell, Bertrand. 1912. *The Problems of Philosophy.* (Oxford, UK: Oxford University Press).
Schnee, Ian. 2015. There is no knowledge from falsehood. *Episteme*, vol. 12, no. 1, pp. 53–74.
Shope, Robert. 1983. *An analysis of knowing.* (Princeton, NJ: Princeton University Press).
Sosa, Ernest. 2007. *A virtue epistemology.* (Oxford, UK: Oxford University Press).
———. 2011. *Knowing full well.* (Princeton, NJ: Princeton University Press).
Unger, Peter. 1975. *Ignorance: A Case for Scepticism.* (Oxford, UK: Oxford University Press).
Warfield, Ted. 2005. Knowledge from Falsehood. *Philosophical Perspectives*, vol. 19, no. 1, pp. 405–416.
Weiner, Matthew. 2006. Must We Know What We Say? *Philosophical Review*, vol. 114, no. 2, pp. 227–251.
Williamson, Timothy. 2000. *Knowledge and Its Limits.* (Oxford, UK: Oxford University Press).
———. 2007. *The philosophy of philosophy.* (Oxford, UK: Oxford University Press).
———. 2011. "Knowledge First Epistemology," in *The Routledge Companion to Epistemology*, ed. Sven Bernecker & Duncan Pritchard, pp. 208–218. (New York, NY: Routledge).
———. 2015. "Justifications, Excuses, and Sceptical Scenarios," in *The New Evil Demon, ed.* Julien Dutant, and Fabian Dorsch. (Oxford, UK: Oxford University Press).