CrossMark

Episteme, 15, 2 (2018) 166–182 © Cambridge University Press
doi:10.1017/epi.2016.53

A DILEMMA FOR THE KNOWLEDGE DESPITE FALSEHOOD STRATEGY

CHRISTOPHER BUFORD AND CHRISTOPHER MICHAEL CLOOS
cb72@uakron.edu, cloos@umail.ucsb.edu

ABSTRACT

One strategy for dealing with apparent cases of knowledge from falsehood is to deny that the knowledge actually is *from* a falsehood. Those endorsing such a move have suggested that cases of knowledge from falsehood are instead cases of knowledge *despite* falsehood. We here provide a dilemma for those wanting to reject the possibility of knowledge from falsehood. The dilemma is explained in part by examining recent attempts to deny that knowledge can be inferentially derived from falsehood.

1. INTRODUCTION

Before introducing the dilemma, we introduce the principles that have been called into question by apparent instances of knowledge from falsehood, as well as a possible counterexample to such principles. Here are three examples of principles that preclude the existence of knowledge from falsehood.

(CCK) If (i) S knows that q, and (ii) believes q solely on the basis of competent deduction from some premises including p, then (iii) S knows that p.

(KP) Inferential knowledge of a conclusion requires known relevant premises.

(TP) Inferential knowledge of a conclusion requires true relevant premises.[1]

And here is one possible counterexample.

HANDOUT
Counting with some care the number of people at my talk, I reason: 'There are 53 people at my talk; therefore my 100 handout copies are sufficient'. My premise is false. There are 52 people in attendance – I double counted one person who changed seats during the count. And yet I know my conclusion. (Warfield 2005: 407–8)

The problem posed by HANDOUT for (CCK), (KP), and (TP) is that Warfield appears to know that he has a sufficient number of handouts even though he has reasoned to that

1 (CCK) is found in Ball and Bloome-Tillman (2014). A variant of (CCK) can also be found in Luzzi (2010). Both (KP) and (TP) are discussed and defended in Montminy (2014). (KP) and (TP) appear to have been first stated by Warfield (2005).

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

conclusion by employing a false belief. That the belief is false rules out (TP); and, if knowledge entails truth, (CCK) and (KP) cannot be correct either.

2. THE KNOWLEDGE DESPITE FALSEHOOD DILEMMA

As noted above, some epistemologists have attempted to defend principles such as (CCK), (KP), and (TP) by arguing that once cases such as HANDOUT are properly understood, it is clear that the knowledge in question is not genuine knowledge from falsehood, where this is understood to require that the falsehood plays an important role in the production of knowledge, but instead is better characterized as *knowledge despite falsehood*. One variant of the knowledge despite falsehood strategy is to argue that what looks to be inferential knowledge obtained from a falsehood is instead inferential knowledge derived from a known proposition. This is the strategy employed by Ball and Bloome-Tillman (henceforth 'BBT'). BBT help illustrate their view by applying it to HANDOUT.

> [W]e think that what is really going on in Handout is something like the following: Warfield knows (perhaps tacitly) that the result of his count is '53'. He then forms two further beliefs: one is the explicit belief that there are 53 people in the room; the other is the tacit belief that there are approximately 53 people in the room. He then employs the latter belief, together with his knowledge that 53 is far less than 100, to conclude that his 100 handouts are sufficient. (BBT: 555)

BBT claim that in cases such as HANDOUT in which it initially appears that a reasoner has inferred a known q from a false p, there will be another proposition t, that is (1) known, (2) entailed by p, and most importantly (3) the real basis of the reasoner's inferential knowledge that q. We will return shortly to an evaluation of BBT's analysis of HANDOUT, but first the advertised dilemma.

We assume that those who employ the knowledge despite falsehood strategy allow for knowledge in scenarios such as HANDOUT that aim to give inferential principles such as (CCK) trouble. Thus defenders of the strategy are tasked with explaining how the beliefs in question, such as Warfield's belief that he has enough handouts, do in fact count as *knowledge* despite the presence of falsehood. The problem is that such defenders agree that the knowledge in question does not come from nowhere; it must come from somewhere. We are dealing with apparent cases of *inferential* knowledge. Those who deny that there are genuine cases of knowledge from falsehood will presumably argue that there exist true propositions that ultimately serve to explain the inferential knowledge in question. However, the role of the falsehood must still be addressed. Might it not still play an important epistemic role? It is for this reason that the dilemma arises. Here is the *knowledge despite falsehood dilemma* explicitly stated.

(D1) Either the falsehood is inferentially inert or it is not.

(D2) If the falsehood is inferentially inert, then cases of knowledge despite falsehood always involve a true proposition that either is implicitly (i.e. tacitly) believed and is

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

itself inferentially inert or is explicitly (i.e. dispositionally or occurrently) believed and is thus an ad hoc amendment to the original cases.[2]

(D3) If the falsehood is not inferentially inert, then there are cases where the falsehood plays an essential role in the explanation of inferential knowledge.

(D4) So, either a true proposition is implicitly believed and is inferentially inert or is explicitly believed and is an ad hoc amendment to the cases, or the falsehood plays an essential role in the explanation of inferential knowledge.

We briefly explain the dilemma. (D1) indicates that the falsehood either plays a role in the inference or it does not. When it does not, it is *inferentially inert*. Inferential inertness is not mere causal inertness. A falsehood could be causally active in producing a conclusion, but that falsehood could fail to be inferentially active in producing the conclusion because the thinker does not accept the conclusion *because* she takes the falsehood to support the conclusion. An inferentially inert proposition will fail to satisfy the Taking Condition.

Taking Condition: Inferring necessarily involves the thinker *taking* his premises to support his conclusion and drawing his conclusion *because* of that fact.[3]

The Taking Condition on inference is supported by the idea that inferential reasoning is person-level and agential. This contrasts with causal processes that are sub-personal and non-agential because

2 We do not intend to take a controversial stand on the metaphysics of belief. Premise (D2) simply considers all relevant possibilities for how the true proposition might be believed by the subject. We follow Eric Schwitzgebel (2015) in drawing the distinction between implicit belief and explicit belief in terms of whether the content of the proposition is represented in one's mind. Schwitzgebel cautions philosophers from confusing implicit belief with dispositional belief, as explicit belief can be dispositional or occurrent. We heed BBT's claim that the true proposition is implicit, tacit, and subconscious. However, (D2) also considers the possibility that the belief is explicitly yet dispositionally held. If one holds an alternative metaphysics of belief, (D2) could be recast in those terms. See Schwitzgebel (2015).

3 From Boghossian (2014). There is debate whether the Taking Condition is as a desiderata for inference. However most theorists in the debate over the nature of inference regard the condition as true. For more on this point see Hlobil (2014). Crispin Wright (2014) is the main opponent of the Taking Condition. Wright thinks even non-reflective, sub-personal processes can count as instances of inference or rule-following. He thinks what is crucial for inference is moving from one information state to another. One need not "take it that" one information state is supported by another in order for thinking to count as inferring. Needing to register that one takes a conclusion to be true because one takes premises to be true is too reflective of a process. Yet, Wright's alternative account of inference is also problematic, as his view contains unanalysed and explained phenomena, and it cannot account for a desiderata of inference called "Inferential Moorean Phenomenon," which states that, "It is either impossible or seriously irrational to infer P from Q and to judge, at the same time, that the inference from Q to P is not a good inference" (Hlobil 2014: 420). For more on problems with Wright's view of inference refer to Hlobil (2014: sect. 3.3). For this paper, we need not fully settle the debate over the nature of inference. Rather, it is enough to embrace an intuitively true principle such as the Taking Condition and note that the sorts of cases involved in the knowledge from falsehood literature involve reasoning that is person-level and attributable to agents as things they can be held responsible for. That is, the Taking Condition best captures the fact that inference is agential and not just something that happens to agents as activities of sub-personal processes. We thank a referee of this journal for urging us to discuss the different perspectives on the nature of inference.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

> no causal process counts as inference, unless it consists in an attempt to arrive at a belief by figuring out what, in some suitably broad sense, is supported by other things one believes. In the relevant sense, reasoning is something we *do*, not just something that happens to us. And it is something *we* do, not just something that is done by sub-personal bits of us. And it is something that we do with an *aim* – that of figuring out what follows or is supported by other things one believes. It's hard to see how to respect these features of reasoning without something like the Taking Condition. (Boghossian 2014: 5)

Inference is a reasoning process and not merely a causal process of transitions between beliefs states or acceptances. Given these considerations, a proposition p is inferentially inert with respect to a particular inference from some premises, which include p, to a conclusion q by a subject S when S does not accept q *because* she takes p to support q. Thus, premise (D1) in the dilemma indicates that either the falsehood plays a role in producing *inferential* knowledge or it does not.

Premise (D2) considers the possibility that the falsehood does not play a role in the inference. If this is the case, the proponent of the knowledge despite falsehood strategy must posit a known proposition responsible for the inferential knowledge. However, if that known proposition is implicitly believed, then it is incapable of playing a role in inference, as the mind does not possess the belief in a way that it can be readily accessed or deployed in inference. Yet, if that known proposition is explicitly believed, then the belief is added to the details of the original case without motivation independent from saving principles (CCK), (KP), and (TP) from being falsified. (D3) considers the possibility that the falsehood is not inferentially inert. Such a possibility results in additional cases of knowledge from falsehood. We show how this occurs in section 4 of our paper.

In the next section, we defend key premises in the dilemma against attempts to falsify them. The knowledge despite falsehood strategy employed by BBT can be understood as an attempt to falsify premise (D2) in the dilemma. For BBT there is always a true proposition that serves as the inferential basis of knowledge, and such a true proposition is not inferentially inert and is not an ad hoc amendment to the cases. We argue that BBT's strategy fails. We then argue that Montminy's defense of the knowledge despite falsehood strategy, involving a rejection of (D3), is unsuccessful as well. Let us now explain how BBT attempt to overcome the dilemma.

3. A DEFENSE OF PREMISE (D2) IN THE DILEMMA

As noted above, BBT argue that cases of knowledge despite falsehood involve a proposition that is distinct from the falsehood. This distinct proposition is the actual inferential basis for the inferred belief that amounts to knowledge. For example, in HANDOUT, BBT claim that Warfield's belief that he has enough handouts is actually inferred from the belief that there are *approximately* 53 people in the room. BBT attempt to undermine the dilemma by identifying a known proposition t that generates and serves as the basis for the belief that p. This amounts to denying premise (D2). The falsehood is inferentially inert, yet there is a true proposition implicitly believed that is not inferentially inert and serves as the basis for the inferential knowledge.

One problem with this proposal concerns the nature of the subject's belief that t. In HANDOUT, BBT describe the subject's belief that there are *approximately* 53 people

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

in the room (here the value of *t*) as an *implicit* or *tacit* belief. BBT do not characterize the tacit belief in much detail, and they do not defend the idea that *t* is tacitly known by way of an account of tacit knowledge.[4] This deficiency in their argument opens them to an objection that if the true proposition is implicitly believed, then it is inferentially inert. They acknowledge this possibility when they say that, "it might be furthermore considered controversial that a subject can have tacit knowledge that plays an active role in his or her reasoning, as required by our account [of knowledge despite falsehood]" (BBT 558). Why is such tacit knowledge inferentially inert?

Inferential reasoning is an agential process. It is not a sub-personal process. Recall, the Taking Condition on inference discussed in the previous section. That intuitively true general condition on inference seeks to capture, "reasoning that is person-level, conscious and voluntary, not sub-personal, sub-conscious, and automatic" (Boghossian 2014: 3).[5] Making an inference is an *act* of a rational agent. Beliefs that are tacit and subconscious cannot play the role in inference of being things the agent takes to support his conclusion.[6] This is problematic for BBT's account because they claim that the true proposition that serves as the actual basis of the reasoner's inferential knowledge is an implicit, tacit, and subconscious belief.[7] Thus, as (D2) indicates, a true proposition that is implicitly (i.e. tacitly) believed is itself inferentially inert.

One way of responding to the argument above is to claim that the true proposition is not implicitly believed but is, rather, explicitly believed. According to this move, the content of the true proposition is represented in the mind of the reasoner, and the proposition is dispositionally or occurrently believed. However, this move runs into the second disjunct in the consequent of (D2) – the move is ad hoc and not well-motivated.

4 As BBT admit, "we will not attempt to justify the claim that [*t*] is known tacitly in this paper by means of the provision of an analysis of tacit knowledge: we take it as given that there is such a phenomenon" (BBT: 558).

5 One might worry that the phenomenon Boghossian seeks to capture is too narrow, as there can be sustaining causes that play a role in inference even though they are not the product of inferences that are conscious. Here it is important to distinguish between beliefs as sustaining causes that are *accessible* to conscious activity and beliefs as sustaining causes that are implicit, tacit, and subconscious such that they are *inaccessible* to conscious activity. We grant that sustaining causes may play a role in inference even though they are not the *product* of conscious inference. What matters is whether such processes involve epistemic bases that the agent takes to support her conclusion. Such "taking" need not be a cognitive process the subject is aware of, though this can be the case; rather, it must be a process that is attributable to the agent or capable of being traced back to, "an attempt to arrive at a belief by figuring out what, in some suitably broad sense, is supported by other things one believes" (Boghossian 2014: 5). Beliefs that are inaccessible to conscious thinking cannot be things the agent takes to support her conclusion. Such sustaining causes will be merely causal in the inferential chain; they will not play an epistemic role in generating the inference. They will be inferentially inert in our sense. We are grateful to an anonymous referee for raising this potential worry.

6 This is because, "One believes *P implicitly* (or *tacitly*) if one believes *P*, but the mind does not possess, in a belief-like way, a representation with that content" (Schwitzgebel 2015). If the mind does not have the belief in a way that it can be readily accessed or deployed in inference, then the person cannot use that belief as something she takes to support her conclusion.

7 As BBT (557 n.9) say about HANDOUT, "we fill in the implicit detail that Warfield holds the mentioned implicit belief and uses it as a base for his knowledge that he has enough handouts." Additionally, as BBT confidently explain, "in the case of *Handout* there are almost certainly some *further*, albeit *tacit*, or *subconscious* beliefs present, which enable him to undertake the inference" (554).

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

BBT amend apparent cases of knowledge from falsehood to serve the purpose of saving principles like (CCK) from refutation. As they claim about HANDOUT, "We think our intuition that the subject of HANDOUT knows his conclusion is tracking certain features of that case which are not spelled out explicitly, but which are nonetheless projected by readers as present" (BBT: 554). It is unlikely that readers are projecting into HANDOUT that Warfield possesses the explicit, occurrent belief that there are approximately 53 people in the room. Such a projection by readers would be reading facts into the case that are not present. It would be an obviously ad hoc way of reading facts into the case in order to argue against knowledge from falsehood. HANDOUT does not indicate that Warfield occurrently believes that there are approximately 53 people in the room. So, the most viable way of articulating the move to explicit belief is that Warfield *dispositionally* believes that there are approximately 53 people in the room. Is this move also ad hoc or is it well-motivated by facts of the case and independent considerations?

The claim that Warfield in HANDOUT believes dispositionally that there are approximately 53 people in the room would be well-motivated given the facts of the case if we had reason to think that Warfield doubts the reliability of his count of the people in the room or it is apparent that we are in a skeptical scenario whereby normal belief-forming processes might be radically misguided. Instead, we are just told that Warfield counts carefully.[8] We are not told that Warfield noticed motion in the room suggesting that people switched seats and he should doubt that he got things right. Additionally, we are not told that this situation is one where philosophical skepticism is salient, as might be the case were Warfield giving a talk on Cartesian skepticism. It is not clear that the context is such that the standards for knowledge ascription are inflated. Were any of the foregoing details explicitly mentioned in the case or readily implied from the facts of the case it would be reasonable to think that readers of the case are attributing a dispositional belief to Warfield about how many people are in the room give or take a few. It seems BBT are reading into the minds of the readers of the case beliefs they take the subject of the case to possess. However these beliefs do not emerge out of the facts of the case; instead, these beliefs are read into the case for the purpose of saving principles like (CCK), (KP), and (TP) from a counterexample.[9] Thus, as (D2) urges, if the true proposition is an explicitly possessed belief, the belief is read into the case and is an ad hoc and unmotivated response.

Despite the foregoing, BBT do attempt to motivate the move to a true proposition serving as the inferential basis for knowledge of the conclusion. They present an altered version of HANDOUT and explain how it supports their position.

8 Recall also that in other suggested cases of knowledge from falsehood, such as FANCY WATCH, where a subject forms a false belief about the time but infers correctly that he isn't late for his meeting, the reasoner possesses an explicit belief in the reliability of his belief-forming mechanism (e.g. Warfield has "extreme confidence in the accuracy of [his] fancy watch"; Warfield 2005: 407). Warfield believes that looking at his watch is a reliable way of forming beliefs about the time. Likewise, it is reasonable to assume in HANDOUT that Warfield explicitly believes that he is reliable at counting people in a room and forming beliefs on that basis. We have no reason to suspect that Warfield has a reason to doubt the reliability of his belief-forming method at the time of the production of his conclusion through inference.

9 One might raise the concern here that without at least a dispositional belief that there are approximately 53 people in the room, Warfield cannot be said to *know* that he has enough handouts. We discuss this objection, offered explicitly in Schnee (2015: section 5).

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

HANDOUT*
Assume that Warfield . . . is informed by someone who just counted the people in the room fifteen times and each time came to the same correct result (52), thus informing Warfield that the number of people in the room is certainly not (exactly) 53. Would Warfield get worried about whether or not he has enough handouts and retract his belief? Surely not. The explanation of why he doesn't, however, is that his belief that he has enough handouts is in fact based on his belief that there are *approximately* 53 people in the room. It is this belief that both causes and sustains Warfield's knowledge; it is therefore this belief upon which Warfield's knowledge is based. (BBT: 556)

HANDOUT* doesn't support the claim that Warfield's dispositional belief that there are approximately 53 people in the room caused and served as the evidential basis for his belief that his 100 handouts are enough. When Warfield is told by his obsessive recounting confidant that there are actually 52 people in the room it is reasonable to think, at that time, Warfield forms the explicit belief that there are approximately the number of people in his original count actually present in the room. It is also reasonable to think that he couples that belief with the belief that 100 is much greater than 52 and reaches his conclusion again that 100 handouts are sufficient. It is fairly common to go through such a quick reevaluation process in light of new information. However this doesn't demonstrate that Warfield's original belief that his handouts are sufficient was based on and sustained by a belief about approximately how many people are in the room. Thus the move to dispositional belief is not antecedently motivated. It results in an ad hoc move involving superimposing details on the cases that are not readily identifiable in the cases or reasonably drawn out of them. Further, this defense of premise (D2) highlights a general problem with the knowledge from falsehood strategy.

The suggestion that there will always be a known proposition that serves as the actual basis for the inferred belief is problematic. It is not enough to always be able to identify a known proposition that evidentially supports the inferred belief. Instead, in order for the account to answer the dilemma, it must always be the case that there exists a distinct known proposition that serves as the *actual* inferential basis for the known inferred proposition. Determining whether the subject's belief that p is caused by the falsehood q or is inferred from the known proposition t appears to be an almost impossible epistemic task, especially given the possibility of subconscious or tacit beliefs.[10]

BBT do provide a case in which they allow that a false belief serves as the basis for an inferred belief.

TEMPERATURE
Bill lives in Pittsburgh, but has to go on a business trip to Toronto in March. He has heard that Canada has long, cold winters, and wonders whether he will need to bring his winter coat. He searches on Google for 'temperature Toronto March' and finds that the average low is −5 degrees.

10 This need for the knowledge despite falsehood strategy to discern the *actual* inferential basis also appears difficult when rendering the known true proposition as dispositionally believed. Of course, I have many dispositional beliefs that might serve as the basis for inference (e.g. that 2674 is greater than 1456), but once this belief serves as the basis of an inference, it appears no longer dispositional. The worry is especially pressing if the subject merely has a disposition to believe t, as opposed to having a dispositional belief that t. For more on this distinction see Audi (1994).

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

> But then he remembers that Canada uses degrees Celsius and wonders how much that is in Fahrenheit. He decides to calculate. Unfortunately, Bill misremembers the conversion formula: although he should add 32 (to get 27), then divide by 9 (to get 3), and finally multiply by 5 (to get 15), he thinks he is supposed to divide by 5 and multiply by 9. At the same time, however, Bill makes typing errors on his calculator: he keys in 5 (rather than −5) then adds 3 and adds 2, rather than adding 32; he divides the result (10) by 5 (to get 2), then multiplies by 9 (to get 18). As a result he believes that the average low in Toronto in March is 18 degrees Fahrenheit, and he concludes he will need his winter coat. (BBT: 556)

If Bill had done the conversion correctly, as articulated by BBT, he would have arrived at the true belief that the average low is 15° in Toronto in March.[11] Instead Bill arrived at the false belief that the average low is 18° in Toronto in March, and he used that false belief to infer the true belief that he will need to bring his winter coat on the trip. BBT think it intuitive that Bill does not know that he needs to bring his coat. Given the process Bill went through, he could have easily not made typing errors and arrived at the belief that it will be 49° in Toronto. This would have led him to infer the false belief that he does not need to bring his winter coat. Thus, the process that yields Bill's belief is not reliable; Bill could have quite easily arrived at a false belief.

Yet, for BBT, it isn't the lack of reliability that causes Bill's belief to not constitute knowledge. Rather BBT (557) think Bill's belief doesn't amount to knowledge because "Bill's explicit numerical beliefs play a crucial role in causing his belief that he won't need his winter coat," and "his false belief that it will be 18° in Toronto in March is not merely epiphenomenal in the causal derivation of his belief that he won't need his coat." BBT posit that it is facts concerning the causal origin of the final belief (i.e. the falsity of the base belief) that best explain why the inferred belief in this case does not amount to knowledge.

The defense offered by BBT is suspect. To see this, notice that there is another plausible explanation for why Bill's final belief fails to amount to knowledge, an explanation that even BBT notice; Bill might easily have had a false belief. In fact, Bill's belief bears all the hallmarks of a classic Gettier case. Linda Zagzebski has identified a common component of most Gettier cases, bad luck is cancelled by good luck (Zagzebski 1994). For example, suppose I receive a postcard from my friend Anand stating he is in Madrid for the summer and form the belief that Anand is in a European capital. However, imagine that Anand has had to rush off to another country to comfort an ailing relative (bad luck). But further suppose that the ailing relative happens to live in Berlin (good luck). My belief that Anand is in a European capital thus remains true though it is not a belief that amounts to knowledge. Bill's belief-forming process is likewise imbued with a combination of bad and good luck. To wit, Bill misremembers the formula (bad luck). But typing errors correct for this mistake (good luck).

11 The correct conversion value should be 23° Fahrenheit. The correct order of operations should have had Bill multiply −5 by 9 (to get −45), divide by 5 (to get −9) and then add 32 (to get 23). BBT presented the conversion order of operations such that Bill should have ended up getting 15° Fahrenheit. However, the difference from 15° Fahrenheit to 23° Fahrenheit does not alter the point BBT make by way of the case. We thank an editor of this journal for bringing this error to our attention.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

If we can eliminate the lack of reliability and the luck component of the example, perhaps we can generate a case of knowledge from false belief. Consider the following variant of TEMPERATURE.

> HEAT ENERGY
> Sal lives in Cleveland, but has to go on a business trip to Toronto in March. He has heard that Canada has long, cold winters, and wonders whether he will need to bring his winter coat. He searches on Google for 'temperature Toronto March' and finds that the average low is −5°. But then he remembers that Canada uses degrees Celsius and wonders how much that is in Fahrenheit. He decides to calculate. Sal correctly remembers the conversion formula and makes no typing errors. However, a small portion of his computer screen was temporarily smudged and as a result he initially believes that the average low in Toronto in March is −8°. Thus, after his calculations he arrives at the belief that the average low is 18° Fahrenheit, and he concludes he will need his winter coat.

In HEAT ENERGY, Sal's explicit numerical beliefs are no less causally relevant than Bill's numerical beliefs in TEMPERATURE. Thus, Sal's false belief that the average low is 18° Fahrenheit is not merely epiphenomenal. Notice though in HEAT ENERGY, Sal's belief that he will need a coat appears safe. The smudge on his screen was minor and his calculations were competent; his belief-formation process was not in danger of leading him to be radically mistaken about the average temperature. Besides being safe, the belief also arguably amounts to knowledge. This though would amount to a refutation of (CCK), (KP), and (TP).

4. A DEFENSE OF PREMISE (D3) IN THE DILEMMA

Martin Montminy (2014) also defends a version of the knowledge despite falsehood strategy. Montminy initially claims to be defending (KP) and (TP).

> Advocates of 'knowledge from falsehood' (KFF) imagine several examples in which a false belief seems to epistemize a subject's knowledge. I will defend both (1) [KP] and (2) [TP] against these alleged counterexamples. (Montminy 2014: 463)

Montminy argues that in all cases of apparent knowledge from falsehood, the falsehood is not essential to the knowledge in question. Montminy is specific about what it is for a belief to be *essential* in this context.

> A belief that p is essential in producing a subject S's knowledge that q just in case S knows that q but would no longer do so if the belief that p were simply removed from S's set of beliefs. (Montminy 2014: 467)

Thus, in HANDOUT, the false belief that there are 53 people in the room is deemed *inessential*. This is because if we were to eliminate just that belief, Warfield would, according to Montminy, still count as knowing that he has enough handouts.[12]

12 Montminy is careful to distinguish his preferred method from the counterfactual method favored by Coffman (2008). For criticism of Coffman's method see Fitelson (2010).

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

One might wonder exactly how the appeal to the inessentiality of Warfield's false belief manages to save (KP) and (TP). Montminy allows that the inferred belief amounts to knowledge and that the base belief is in fact false. Isn't this sufficient to show that both (KP) and (TP) are false? They are if we understand inferential knowledge that *q* as involving a belief that is arrived at (a) via inference from a belief that *p* and (b) amounts to knowledge. However, if we understand inferential knowledge that *q* as involving not only (a) and (b), but also a third condition (c) *q* amounts to knowledge *because* it was inferred from *p*, then perhaps (KP) and (TP) can be saved. In cases such as HANDOUT, that condition (c) obtains is not clear. Montminy adds,

> I seriously doubt that we intuit that it is S's inferring q from the false belief that p that grants S knowledge that q. Our intuitions are less specific: in each case, S infers q from p and S counts as knowing that q. Hence intuition alone does not support the claim that S acquires knowledge that q by inferring q from p. (Montminy 2014: 466)

Returning to our dilemma may help to illuminate Montminy's position. Montminy appears to reject (D3). He argues that even though the knowledge that *q* might come from *p* via inference (thus *p* is not inferentially inert) this does not falsify principles such as (KP) or (TP) since such principles concern inferential knowledge in the strict sense (i.e. as requiring (a–c)). On Montminy's view, one does not know *q* *because* it was inferred from falsehood *p*, as one can remove *p* from one's evidential set and continue to know *q*.

Montminy's response to the dilemma thus hinges on whether the false belief in apparent cases of knowledge from falsehood really is inessential; it is not obvious that this is so. Consider the following case, which is a variation of HANDOUT.

> TEST
> Martin counts carefully each person in the exam room. On the basis of this method, he forms the belief that there are exactly 53 students in the room. The belief is false since one person switched chairs while he was counting; there are only 52 students in the room. He infers that his 100 test copies are sufficient. Right after forming this belief, he runs into a graduate student, a notoriously unreliable estimator of room occupancy, who tells him that he has plenty of handouts.

In TEST, Martin knows that he has enough copies of the test. The key question is whether the false belief that there are exactly 53 people in the room is essential to his knowledge that he has enough copies. Montminy's position requires that it not be essential since, if it was, we would have a genuine case of knowledge from falsehood. Given Montminy's definition of essential, the false belief in TEST is essential in producing his knowledge that he has enough copies if it's true that were his belief that there are exactly 53 people simply removed from his set of beliefs, he would no longer know that he has enough copies.

Suppose then that we consider Martin with his current set of beliefs minus the belief that there are exactly 53 persons in the room. Does he know that he has enough copies of the test? It is certainly true that he possesses evidence that supports this belief. However, possessing evidence that supports a belief and having that belief does not entail knowledge since one might have formed the belief for some other reason. And the other reason might be faulty.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

When we remove the belief that there are exactly 53 people in the room from Martin's set of beliefs, we are left with no answers as to how Martin actually formed his belief that he has enough copies. It is true that had Martin not employed the false belief, he could have relied on the truth that there are approximately 53 people in the room.[13] Yet if it is this counterfactual fact which rules out the essentiality of the false belief, then Montminy's account collapses into Coffman's (2008) account, which holds the truth of (6),

(6) If S's inferential belief that q had been based on the belief that p' [the approximate truth that follows from p], the inferential belief would (still) have constituted knowledge,

to be sufficient to establish that the false belief that p is not an essential falsehood. Further, Fitelson (2010) has argued persuasively, and successfully according to Montminy (2014: 468), that (6) can be false even though in the actual world a falsehood appears to serve as the inferential basis for knowledge.

We can press the worry with Montminy's test for essentiality if we stipulate that though Martin forms his belief that he has enough handouts on the basis of his false belief that there are exactly 53 people in the room, had he not done so, he would have relied on his graduate student's testimony. Martin does know that he has enough handouts, but in this new scenario, removing the falsehood from Martin's set only weakens the case for claiming that he still knows, on the basis of his dispositional belief that there are approximately 53 people in the room, that he has enough handouts. Thus, as premise (D3) articulates, if the falsehood is not inferentially inert, there will be cases of knowledge from falsehood.

5. A DEFENSE OF ESSENTIAL KNOWLEDGE FROM FALSEHOOD

Ian Schnee (2015) endorses an alternative understanding of essential falsehood. He uses that understanding to argue that there are no cases of knowledge from falsehood. In this section, we defend the possibility of knowledge from falsehood from Schnee's dilemma. In order to fully grasp the dilemma, it will help to explicate a few key distinctions introduced by Schnee. First, there is the distinction between *S-evidence* and *K-evidence*.

> I distinguish between a piece of evidence that a subject uses in a conscious inference and the evidential basis of her knowledge (the grounds on which she has inferential knowledge); call the first notion **S-evidence** and the second notion **K-evidence**. (Schnee 2015: 54)

This distinction allows us to make Montminy's earlier point in a more perspicuous fashion; philosophically interesting cases of knowledge from falsehood should involve falsehoods that function not only as S-evidence but also as K-evidence. Second, Schnee uses this distinction to draw a line between knowledge from falsehood *in general* and *essential* knowledge from falsehood.

13 Thank you to an anonymous referee for raising this possibility.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

> I will therefore understand KFF in general to be any cases in which a subject has knowledge grounded on false K-evidence, and essential KFF to be a case of KFF in which the falsehood is part of the sole evidential path enabling the subject's knowledge. (Schnee 2015: 56)

Schnee argues that there is not essential knowledge from falsehood and that if there is no essential knowledge from falsehood, then there is no knowledge from falsehood in general. In this section, we examine his argument against essential knowledge from falsehood. We address the conditional claim in the next section.

Schnee begins by examining potential cases of knowledge from falsehood. Here is his first example.

> TV SHOW 1: Ellie's favorite TV show is on from six to seven. Ellie looks at her watch, which reads six thirty, and she infers, from the consideration that it is (exactly) six thirty, that her show is on. But Ellie is in error. It isn't (exactly) six thirty, it is six thirty-two; her watch is slow by two minutes. (Schnee 2015: 54)

Schnee notes that, in TV SHOW 1, knowledge is consistent with it being the case that Ellie's belief that is approximately six thirty functions as K-evidence. Given the assumption that essential knowledge from falsehood must involve one and only one evidential knowledge enabling path, what is needed is a case of knowledge from falsehood where it is explicitly stated that the subject does not believe any truths that are 'approximate' versions of the falsehood in question. Thus, Schnee points to the following example.

> TV SHOW 2: This case is like TV Show 1, but Ellie is extremely confident about the accuracy of her watch and only forms exact beliefs from it; she has no other evidence at all regarding the time. (Schnee 2015: 55)

With TV SHOW 2 in hand, we can now turn to Schnee's argument against essential knowledge from falsehood. Schnee's argument can also be presented in dilemma form.

(1) Either Ellie has a good epistemic reason for not forming a belief about approximately what time it is or she does not have such a reason.

(2) If Ellie has a good reason for not forming beliefs about approximately what time it is, then Ellie is Gettiered and lacks inferential knowledge.

(3) If Ellie does not have a good reason for not forming beliefs about approximately what time it is, then Ellie is epistemically defective and lacks inferential knowledge.

(4) Thus, Ellie lacks inferential knowledge.

Schnee defends (2) by presenting a series of cases, starting with TV SHOW 3, which involve Ellie having good reason for not forming a belief about approximately what time it is. In these cases, Schnee argues that even if Ellie's inferred belief is true, it doesn't amount to knowledge.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

> TV SHOW 3: This case is like TV Show 2, but Ellie knows that, even though her watch is quite reliable, when her watch is not exactly right it is usually an hour or more off. That is why she does not form approximate beliefs about the time from her watch; she believes exactly what her watch says or she forms no belief at all about the time on its basis. (Schnee 2015: 58)

Schnee seems correct that if Ellie does infer from the false belief that it is six thirty the true belief that her show is still on, she does not acquire knowledge. This is due to the fact that she believes (and knows) that if her watch is off, then it is usually quite wrong. However, Ellie might have other reasons for not forming approximate beliefs. Consider such a case.

TV SHOW: This case is like TV Show 2, but Ellie does not form approximate beliefs about the time from her watch because she purchased her watch from a watchmaker who has made 10,000,000 watches and only one has been less than fully accurate. And that watch was only off by one second. Her watch does not however track seconds, only the hour and minutes. Ellie has checked her watch numerous times and it has always delivered a correct verdict with respect to hour and minutes. She is thus extremely confident that her watch is correct and refuses to form beliefs about the approximate time (perhaps viewing such beliefs as unnecessary clutter). Her watch is in fact off by one second. She forms her belief at 6:29:59.[14]

In TV Show, Ellie appears to have good epistemic reason for forming only exact beliefs about the time; the odds of her watch being incorrect are astonishingly low. However, there is some pull towards saying that she does know that her show is still on. Ellie has had the bad luck of receiving a less than perfect timepiece. But though her watch is in fact incorrect, it is only off by a second and is not a wildly defective watch. It is hard then to see why Ellie being correct about the show still being on is properly described as being the result of the good luck that is typical of Gettier-cases. It also isn't obvious that Ellie's belief fails to be safe. If she forms her belief at roughly six thirty, then in close possible worlds where she forms the belief that her show is on, it is on.[15] There is thus ample reason for thinking that (2) is false. And the falsity of (2) would be sufficient to disarm the dilemma.

There are also concerns that can be raised with premise (3) in Schnee's dilemma. Even if Ellie lacks *good* reasons for refusing to form beliefs about the approximate time, a case can be made that she can still arrive at inferential knowledge from a falsehood. Imagine that Ellie wears a reliable watch from a reliable manufacturer and that the watch happens to be off by a few minutes. Suppose further that Ellie suffers from a strain of obsessive-compulsive disorder that keeps her from forming beliefs about the approximate time; she feels a compulsion to only have fine-grained beliefs about the time. This isn't a good epistemic reason for refusing to form beliefs about the approximate time, but why think that it rules out Ellie coming to know that her show is on?

14 Thanks to an anonymous referee for comments that led to needed revision of the structure of this scenario.

15 Schnee (2015: 62) states of a similar case, "in nearby worlds in which she looks at her watch when it reads 6:59, she will form the belief that her show is on, even though that belief is false . . ." Such worlds may be nearby, but we disagree that they are necessarily relevant to evaluating the safety claim. Objects can be closer than those that are merely nearby.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

6. IMPLICATIONS FOR THE NATURE AND STRUCTURE OF JUSTIFICATION

In this section, we uncover important implications of our defense of knowledge from falsehood and we use those implications to make further points against the knowledge despite falsehood strategy and in favor of the possibility of knowledge from falsehood. One interesting result of our defense of the possibility of knowledge from falsehood is that it prompts the question of how relevant the actual inferential basis of a belief is when a subject possesses evidence beyond the basing belief.

Many epistemologists have endorsed a distinction between *propositional* and *doxastic* justification, which consists in the distinction between having justification to believe a proposition and having a justified belief in that same proposition. An example to help illuminate this distinction might involve a subject who has strong testimonial evidence in favor of a proposition but only believes the proposition due to wishful thinking. Here the subject possesses propositional justification, but does not have a justified belief since the belief is not properly based on the evidence. Yet, some defenders of knowledge despite falsehood, such as Montminy, suggest that in some cases the possession of certain kinds of evidence (and thus certain cases of propositional justification) may entail that a belief is justified (and perhaps known) independently of consideration of the inferential basis of the belief. For example, as discussed above, Montminy allows that in certain cases the actual etiology of the subject's belief is irrelevant since the false belief that serves as the basis for the inferred belief is inessential to the production of knowledge. We have argued that Montminy's view is flawed, but our criticism has led to interesting questions about the nature and structure of justification.

Those who think that knowledge cannot come from falsehood presumably do so on the grounds that beliefs that amount to inferential knowledge cannot be *properly based* on false beliefs. Peter Murphy (2015) has suggested that a false unjustified belief may still serve as the proper basis for a justified belief. While Murphy's discussion is not directly concerned with knowledge from falsehood, his discussion of proper basing involves a case of a false belief serving as the proper basis for an inferred justified belief. Murphy introduces the concept of a *justification conduit* to help explain the possibility of such knowledge.

> The second idea is the novel epistemic status of being a *justification conduit*. This status belongs to some *unjustified beliefs* relative to some beliefs that are validly and competently inferred from them. In particular, it belongs to beliefs that are themselves unjustified but nonetheless allow justification to pass through to some beliefs that are validly and competently deduced from them. (Murphy 2015: 11)

Murphy relies on the following scenario to illustrate the possibility of such a conduit.

> FRED
> Fred is very busy. Each week, he has at least twenty scheduled meetings. Going only on his memory, he believes that he has a meeting with Mary next Wednesday at noon. However, Fred's memories about the exact date and time of his meetings are often mistaken: when he has a meeting scheduled for a Tuesday, Wednesday, or Thursday, his memory is often off by a full day; and his memories about the exact time of his meetings are also frequently wrong. Fred knows these things about himself. But he also knows that he is very good at determining *what week* a scheduled meeting falls on when he infers this from one of his memory-based beliefs about that meeting's exact date and time. In the present case, Fred infers from his memory-based belief that he has a

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

> meeting with Mary at noon next Wednesday to the conclusion-belief that he has a meeting with Mary sometime next week. (Murphy 2015: 6–7)

Murphy also argues that FRED supplies a counterexample to the following principle.

> Single-Premise JFJ: Where S's belief that p is inferred from one premise-belief and is not epistemically overdetermined, S's belief that p is inferentially justified only if that premise-belief is justified.

According to Murphy, one's memory representation may not justify the belief that the meeting is on Wednesday, but, because of the inference, it might serve to justify the belief that the meeting is next week.

The notion of a justification conduit is worth investigating, but it isn't clear that Fred's situation provides us with a case of such a conduit. The clearest case of a justification conduit would involve a single inference from an obviously unjustified premise-belief to an obviously justified conclusion-belief. However, if we pare down the case of Fred to include only facts about the inference from Fred's belief that he has a meeting on Wednesday to Fred's belief that the meeting is next week, it looks as if the premise-belief is not capable of serving as a justification conduit. Either the premise-belief is justified, in which case it can't serve as a justification conduit, or the conclusion-belief is unjustified, in which case conductance has failed. To bolster the claim that the conclusion-belief is justified, Murphy adds that Fred is aware of his unreliable memory-based beliefs about the exact date of his meetings (thus making the premise-belief unjustified) and the fact that when he is wrong about the exact date of the meeting, he is still often right about the week of the meeting (thus allowing the conclusion-belief to be justified). The problem that arises due to this maneuvering is that it is questionable whether the initial evidence for the unjustified belief (i.e. the memory representation) is doing *any* justificatory work in Fred's case. It looks as if Fred's justificatory structure can be captured with the following argument.

(1) I have a memory-based belief that the meeting is on Wednesday.
(2) In the past, when I've had such a belief about the date of the meeting, I've at least been right about the week of the meeting.
(3) Thus, I have a meeting next week.

Of course, the memory representation here is presumably causally responsible for the fact that Fred has a belief about the date of the meeting. However, notice that (1) above is not a premise-belief *based* on such a representation, but it is instead a belief *about* such a belief. Fred's situation thus begins to look like a case of justified belief *despite* rather than *from* unjustified belief.

We propose that rather than serving as justification conduits, inferences that yield cases of knowledge from falsehood serve as justification *tethers*. The inferences in question (e.g. from the false belief that there are exactly 53 people in the room to the belief that one has enough handouts) ensure that the inferred beliefs are safe since similar inferences in nearby possible worlds do not lead to falsehoods. The false beliefs in question are also part of the tethering process since facts about such beliefs, and how they were formed, allow us to assess questions concerning luck and reliability. For example, though Warfield's belief that there are exactly 53 people in the room is false, intuitions about the inferred belief amounting to knowledge would likely differ if Warfield's initial belief was that 99 people

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

were in the room. Thus, the tethering role played by the false belief is essential to the inferential knowledge produced by the falsehood.

The notion of tethering allows us to now address Schnee's conditional claim that if there is no essential knowledge from falsehood, there is no knowledge from falsehood in general. Schnee defends this claim by pointing out that if there are multiple inferential paths to an inferred truth, say one involving a falsehood and one involving a truth, then it will always be possible that the path involving the truth ultimately explains the knowledge in question.

> My argument against KFF requires, for any putative KFF case in which the subject possesses knowledge, that there is an alternate K-evidential path containing only truths that actually enables the subject's knowledge. (Schnee 2015: 67)

This commitment opens Schnee's account to an objection similar to the objection raised earlier against Montminy. To wit, there will be cases where the false belief tethers the inferential knowledge to the evidence. In a case such as TEST, Schnee must claim that the false belief plays no role in enabling knowledge. However, given the existence of multiple evidential paths, one good and one poor, the false belief actually helps properly ground the inferred belief. Without the false belief, the possibility that Martin grounds his belief improperly cannot be ruled out.

7. CONCLUSION

We conclude that although proposed cases of knowledge from falsehood may involve known truths, there is ample room for falsehoods to play an essential role in the production of inferential knowledge. Ball and Bloome-Tillman's defense of principles requiring knowledge to be derived from truths suffers from a failure to show that a known truth can serve as the inferential basis in every possible case of knowledge from falsehood. Further, Montminy's defense rests upon his inadequate definition of *essentiality*. Schnee's arguments are also unsuccessful. There is reason to believe that the lack of an independent evidential path need not preclude knowledge. And the existence of alternative evidential paths does not rule out false beliefs playing an interesting and important epistemic role. We have suggested that understanding false beliefs as playing a tethering role allows us to see the vital role false beliefs might play in the generation of genuine inferential knowledge.[16]

REFERENCES

Audi, R. 1994. 'Dispositional Beliefs and Dispositions to Believe.' *Noûs*, 28: 419–34.
Ball, B. and Bloome-Tillman, M. 2014. 'Counter Closure and Knowledge Despite Falsehood.' *The Philosophical Quarterly*, 64: 552–68.
Boghossian, P. 2014. 'What is Inference?' *Philosophical Studies*, 169: 1–18.

16 Thank you to an anonymous reviewer and the editors of this journal for helpful comments. Thank you as well to John Chaplin for discussion.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press

Coffman, E. J. 2008. 'Warrant Without Truth?' *Synthese*, 162: 173–94.
Fitelson, B. 2010. 'Strengthening the Case for Knowledge from Falsehood.' *Analysis*, 70: 666–9.
Hlobil, U. 2014. 'Against Boghossian, Wright and Broome on Inference.' *Philosophical Studies*, 167: 419–29.
Luzzi, F. 2010. 'Counter-Closure.' *Australasian Journal of Philosophy*, 88: 673–83.
Montminy, M. 2014. 'Knowledge Despite Falsehood.' *Canadian Journal of Philosophy*, 44: 463–75.
Murphy, P. 2015. 'Justified Belief from Unjustified Belief.' *Pacific Philosophical Quarterly*. Advance Online Publication. doi: 10.1111/papq.12129.
Schnee, I. 2015. 'There is No Knowledge From Falsehood.' *Episteme*, 12: 53–74.
Schwitzgebel, E. 2015. Belief. In *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/belief/
Warfield, T. 2005. 'Knowledge from Falsehood.' *Philosophical Perspectives*, 19: 405–16.
Wright, C. 2014. 'Comment on Paul Boghossian, 'What is Inference'.' *Philosophical Studies*, 169: 27–37.
Zagzebski, L. 1994. 'The Inescapability of Gettier Problems.' *The Philosophical Quarterly*, 44: 65–73.

CHRISTOPHER BUFORD is Associate Professor of Instruction at the University of Akron, Ohio. His philosophical interests include epistemology, personal identity, and bioethics.

CHRISTOPHER MICHAEL CLOOS is a candidate in philosophy at the University of California, Santa Barbara. His primary research interests are in epistemology, ethics, and free will and moral responsibility.

https://doi.org/10.1017/epi.2016.53 Published online by Cambridge University Press