



Kelp, C. (2016) Epistemic Frankfurt cases revisited. *American Philosophical Quarterly*, 53(1), pp. 27-37.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/140964/>

Deposited on: 16 May 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Epistemic Frankfurt Cases Revisited

Christoph Kelp*

Abstract

In [Author 2009a] I argued that there are epistemic Frankfurt cases that serve to show that knowledge does not require safety from error. In this paper, I revisit these Frankfurt cases. I first argue that a recent response to my earlier argument by Duncan Pritchard remains unsatisfactory. I then show that Frankfurt cases impact a much wider range of accounts. Specifically, I argue in some detail that, in conjunction with the infamous Fake Barn cases, they generate a problem for the two most prominent virtue theoretic accounts to knowledge, due to Ernest Sosa and John Greco. Finally, I conclude by offering some reason to think that the lesson these cases teach may just be that a knowledge first approach to epistemology is favourable to its traditional rival.

1 Introduction

Since Gettier's seminal [1963] article, one of the key tasks in epistemology has been to identify a condition that, in conjunction with justified true belief, is sufficient for knowledge. Whilst a lot of ink has been spilled on this project, it is fair to say that the search for a successful 'anti-Gettier' condition, and hence for a workable account of knowledge, is ongoing.

That said, there are a number of live proposals on the market. Among the most promising candidates are safety-based [e.g. Sosa 1999, Pritchard 2005, 2012a] and virtue epistemological [e.g. Greco 2010, 2012, Sosa 2010, 2011] accounts of knowledge. The core idea of safety-based accounts is that knowledge requires safety from error.

*Centre for Logic and Analytic Philosophy, KU Leuven, Kardinaal Mercierplein 2, BE-3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

Roughly, this means that knowledge requires one to avoid error at nearby possible worlds. In contrast, virtue epistemology (*VE*) has it that knowledge is a kind of success because of ability. According to standard versions of the view, the relevant kind of success at issue in knowledge is true belief and the relevant kind of ability is cognitive ability, that is, roughly, a disposition to form true beliefs.

In a recent paper [Author 2009a], I argued that no matter whether safety-based accounts of knowledge can successfully handle Gettier cases, they are in any case too strong to be necessary for knowledge. There are cases in which an agent has an unsafe belief that nonetheless intuitively qualifies as knowledge. In this paper, I will return to this argument and show that a recent response to my earlier argument, due to Pritchard, fails (§2). I will then argue that the relevant cases also serve to generate a problem for Greco and Sosa's virtue epistemological accounts of knowledge (§3). Finally, in the conclusion (§4), I will give some reason to think that the lesson to be learned here may just be that we should opt for a knowledge first approach to epistemology.

2 Safety-Based Epistemology

2.1 The Safety Condition on Knowledge

Safety-based epistemology claims that one knows a proposition, p , only if one safely believes that p . (I will henceforth also refer to this thesis as 'the safety condition on knowledge' or simply 'the safety condition' for short.)

While the core of safety is captured in the idea of avoidance of error at nearby worlds, there is reason to believe that, in the final analysis, the safety condition will need a number of further refinements, including at the very least an index to ways of belief formation [e.g. Pritchard 2005, following Nozick 1981]. It is worth noting that, if successful, my argument will work against a variety of ways in which one might implement such refinements. For instance, it will work against strong versions of the safety condition according to which knowledge requires safety from error across all nearby worlds as well as weak versions that require safety from error across most nearby worlds. It will also work against belief-based versions of the safety condition, according to which knowing p requires avoiding false belief that p across (the relevant range of) nearby worlds, as well as method-based versions, according to which knowing p via method

M requires that *M* not produce false beliefs across (the relevant range of) nearby worlds. For the purposes of this paper, however, I will be focusing on Pritchard's account of safety, which he states in the following passage:

S's belief is safe iff in most near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as the actual world, her belief continues to be true.¹

[Pritchard 2007: 292, 2009: 34]

With Pritchard's account of safety thus in play, let's move on to the argument that knowledge does not require safety. At its heart is the following case:

FRANKFURT CLOCK. Russell's arch-nemesis, a powerful demon, has an interest that Russell forms a belief that it's 8:22 by looking at the grandfather clock in the hallway when he comes down the stairs. Russell's arch-nemesis is prepared to do whatever it may take in order to ensure that Russell acquires a belief that it's 8:22 by looking at the grandfather clock when he comes down the stairs. However, Russell's arch-nemesis is also lazy. He will act only if Russell does not come down the stairs at 8:22 of his own accord. Suppose, as it so happens, Russell does come down the stairs at 8:22. Russell's arch-nemesis remains inactive. Russell forms a belief that it's 8:22. It is 8:22. The grandfather clock is working reliably as always.²

Intuitively, Russell knows that it is 8:22 in this case. After all, we may assume, he has the ability to read the clock and forms a belief via an exercise of this ability. Moreover, the clock is actually functioning properly and the reading is accurate. At the same time, Russell's belief is not safe. After all, among the very close nearby worlds are worlds at which Russell comes down the stairs a minute earlier or later. At those worlds, the demon intervenes with the result that Russell takes a reading from an inaccurate clock and ends up with a false belief that it is 8:22. At the same time, Russell continues to form his belief in the same way as in the actual world, i.e. by exercising his clock-reading ability. Since safety does not tolerate false beliefs at very close nearby possible worlds at which the agent forms his belief in the same way as in the actual world, Russell's belief turns out unsafe. The safety condition predicts, incorrectly, that Russell does not know.

2.2 Pritchard's Response

Pritchard responds to this argument by denying that Russell knows that it's 8:22. In support of this claim he offers a couple of arguments and an account of how those who have the intuition could have become entrapped in error. Here is his first argument:

[Russell] is effectively finding out the time by looking at what is (for him anyway) a stopped clock, since whatever time he comes downstairs the clock will say '8.22 a.m.'. But one cannot gain knowledge about the time by consulting a stopped clock, even when one happens to form a true belief!

[Pritchard 2012b: 183, 2013: 160]

I must confess that I don't find this argument particularly convincing. To begin with, notice that it is a crucial part of the case that the clock is functioning properly and is thus *not* stopped. The issue of whether the clock is stopped is no more agent-relative than the issue of whether the earth revolves around the sun. It is no more plausible that, in FRANKFURT CLOCK, the clock is stopped for Russell, than it is plausible that the sun revolves around the earth for Ptolemy. Of course, Pritchard is free to define an agent-relative concept of stopped, but then he evidently cannot appeal to the plausible idea that one cannot acquire knowledge from a stopped clock in order to argue that an agent who satisfies this agent-relative concept does not know.

There is further reason to believe that Pritchard's argument will not do the trick for him. To see this, consider the following variation of the case:

FRANKFURT CLOCK*. Russell's demon nemesis wants Russell to believe either that it is 8:22 or else that it is past 8:22. He has resolved that if Russell comes down before 8:22 he will set the clock to 8:22 and if he comes down at 8:22 or later, he will do nothing. Russell comes down at 8:22, the demon remains inactive, and Russell acquires a true belief by taking a competent reading from a perfectly functioning clock.

The intuition of knowledge remains unaffected here. Moreover, Russell's belief remains unsafe. After all, Russell might so easily have come down the stairs a minute earlier, in which case the demon would have manipulated the clock with the result that the reading would have been inaccurate. Finally, however, in the present version

of the case, the clock is not stopped for Russell in the sense envisaged by Pritchard. To see this, note that it is not the case that whatever time Russell comes downstairs the clock will read '8:22'. Russell might very easily have come down a minute later, in which case the demon would also have remained inactive and the clock would have read '8:23'. If so, Pritchard's response is bound to fail anyway.

Let's move on to Pritchard's second argument that Russell does not know. In order to see how it works, we will first have to look at a different case:

FAKE BARNS. Barnes, a reliable barn spotter, is driving through the countryside. He looks out of the window, sees a barn and comes to believe that he is looking at a barn. Whilst Barnes's belief is true, unbeknownst to him, the structure he is looking at is the only real barn in an area that is otherwise peppered with fake barns that are indistinguishable from real barns, at least from Barnes's position on the road [Goldman 1976: 772-73].

I take it to be agreed that Barnes does not know that he is looking at a barn. Crucially, in Pritchard's view, **FRANKFURT CLOCK** is structurally analogous to **FAKE BARNS**. Since, in **FAKE BARNS**, Barnes lacks knowledge, we should accept that the same goes for Russell in **FRANKFURT CLOCK** [Pritchard 2009: 40].

There are a number of problems with this argument. To begin with, it is not really clear to me that **FRANKFURT CLOCK** and **FAKE BARNS** are indeed analogous. For starters, the very fact that they elicit diverging intuitive verdicts provides some evidence that they are not analogous after all. What's more, even if the two cases are relevantly analogous, it is far from clear that the lesson to be learned is that, in **FRANKFURT CLOCK**, Russell lacks knowledge rather than that, in **FAKE BARNS**, Barnes has knowledge. At the very least, this point affords argument. Unfortunately, Pritchard fails to deliver on this front.

That said, Pritchard might be able to discharge at least part of the burden that he now finds himself settled with by offering a plausible error theory for the intuition that Russell knows. In order to achieve this, Pritchard distinguishes between knowledge on the one hand, and cognitive achievement on the other. A cognitive achievement is, roughly, a cognitive success that is primarily creditable to cognitive ability. Knowledge, in contrast, features both a weaker ability condition—i.e. it requires cognitive success that is not primarily but to a significant degree creditable to cognitive ability—and

the safety condition. The idea then is that, while knowledge and cognitive achievement often go hand-in-hand, they can sometimes come apart. Cases like FRANKFURT CLOCK constitute one type of case in which they do. Those who have the intuition that Russell has knowledge in this case mistake cognitive achievement for knowledge [Pritchard 2009: 40, 2012b: 183, 2013: 160].

There are two problems with Pritchard's error theory. To see the first, notice that what Pritchard owes us at this stage is an explanation of why our intuition leads us astray in FRANKFURT CLOCK *but not* in FAKE BARNS. Pritchard's story does not achieve this. After all, FAKE BARNS, too, is a case of cognitive achievement without knowledge. If, in such cases, we are prone to mistake cognitive achievement for knowledge, the question remains why we don't do so in FAKE BARNS. The issue that affords diagnosis—why intuition leads us astray in FRANKFURT CLOCK but not in FAKE BARNS—thus remains unresolved.

Pritchard's diagnosis fails on yet another and more significant count: there are variations of FRANKFURT CLOCK in which the agent does not secure a cognitive achievement, at least not by Pritchard's own lights. To see this, notice that, besides cases of cognitive achievement without knowledge, Pritchard also grants that there are cases of knowledge without cognitive achievement. Most importantly, cases of testimonial knowledge in which the greatest part of the cognitive work is done by the testifier and not the receiver of testimony fall under this heading [Pritchard 2008: 446 and Pritchard et al. 2010: 40-43]. That means that if we can offer a (suitable) variation of FRANKFURT CLOCK involving testimonial knowledge in which the bulk of the cognitive work is done by the testifier, Pritchard is committed to treating the case as one in which the agent does not secure a cognitive achievement. If so, his diagnosis of FRANKFURT CLOCK in terms of mistaking cognitive achievement for knowledge is bound to fail.

This leaves us with the task of providing a suitable case. Here goes:

FRANKFURT CLOCK**. As Russell comes down the stairs his little cousin, who cannot yet read the clock, asks him what time it is. Russell looks at the clock, sees that it reads 8:22 and tells his cousin that it is 8:22. On that basis, Russell's cousin acquires a true belief that it is 8:22. Unbeknownst to both of them, the cousin's demon nemesis was lurking in the background, prepared to set the clock to 8:22 had Russell not happened to look at it at just the right time, i.e. at precisely 8:22.

Just as in *FRANKFURT CLOCK* before, intuitively, the cousin's belief qualifies as knowledge. At the same time, this belief is not safe. At many nearby possible worlds, including very close ones, at which Russell's cousin acquires the belief via testimony from Russell, Russell comes down the stairs a minute earlier or later. Since, in that case, the demon shows his hand, at those worlds the cousin ends up with a false belief. Crucially, the bulk of the cognitive work is done by the testifier (here: Russell). After all, Russell is the one who has the crucial cognitive ability in this case. If so, by Pritchard's lights, the belief about the time Russell's cousin forms does not qualify as a cognitive achievement. We thus have a case of unsafe knowledge without cognitive achievement. Pritchard's diagnosis of the intuition of knowledge in terms of mistaking cognitive achievement for knowledge fails on yet another count [see also *Author 2009b*]. Despite Pritchard's efforts, then, safety-based accounts of knowledge remain in trouble.

3 Virtue Epistemology

Thus far, I have put Frankfurt cases to use against safety-based accounts of knowledge. In this section, I will try to show that they have a wider-ranging impact as they also serve to generate a problem for the two leading virtue theoretic accounts of knowledge, due to John Greco and Ernest Sosa. More specifically, I will argue that both Greco and Sosa's versions of *VE* have the undesirable consequence that agents in Frankfurt cases lack knowledge as well, at least so long as they hold on to their respective accounts of Fake Barn cases. In order to see how the argument works, I will first have to briefly sketch Greco and Sosa's virtue epistemological accounts of knowledge and their accounts of Fake Barn cases. I will start with Greco.

3.1 Greco

Greco takes abilities to be dispositions to achieve results within a certain range. Moreover, abilities are relative to what Greco calls "environments" (*E*), that is "sets of relatively stable circumstances", and "conditions" (*C*) or "sets of shifting circumstances within an environment" [Greco 2010: 77]. This core idea is unpacked in more detail as follows:

S has an ability $A(R/C)$ relative to environment $E =$ Across the set of relevantly close worlds W where S is in C and in E , S has a high rate of success in achieving R .

[Greco 2010: 77]

In case of the cognitive abilities at issue in his version of VE , $R =$ true belief (in propositions within a certain range). In order to possess the relevant kind of cognitive ability relative to C and E , one must be such that, at close worlds at which one is in C and E , one attains true belief about propositions in the range with a high rate of success.

In addition, Greco maintains that attributions of abilities afford a contextualist semantics: context determines which *ability to R in C and E* is picked out by 'ability to R '. In the case of attributions of the cognitive abilities at issue in VE the story is even more complex: context first determines a practical reasoning context, which may be the agent's, the attributors' or some third party's. The relevant cognitive ability picked out by 'cognitive ability' is then fixed in accordance with what Greco takes to be a core function of the concept of knowledge, to wit, flagging actionable information and sources of information, where what is actionable is determined by the practical reasoning context fixed at the first step [Greco 2010: 78-79].

Finally, Greco unpacks the because relation at issue in VE in terms of explanatory salience: a success is because of ability just in case ability plays a salient or important part in the causal explanation of the success. Crucially, abilities have default salience in such explanations. This default salience may be trumped by abnormal factors, which are salience magnets, as it were [Greco 2010: 73-75].

According to Greco, this is precisely what happens in standard Gettier cases. For instance, when an agent acquires a true belief about the time by taking a reading from a clock that stopped exactly twelve hours earlier [Pritchard 2005: 137-38, which he takes from Russell 1948], the fact that the clock stopped exactly twelve hours earlier constitutes the abnormal factor which trumps the default salience of the agent's ability. What is salient in the causal explanation of why the agent forms a true belief about the time here is not the fact that he exercises his ability to read the clock. Rather, it is the fact that he happens to look at the clock exactly twelve hours after it stopped. For that reason, even though the agent exercises his clock-reading ability and attains cognitive success, his success is not because of the exercise of his ability. In this way, Greco can explain the absence of knowledge in such cases.

Let's now take a closer look at how Greco deals with FAKE BARNS. As a first observation, Greco grants that this case cannot be dealt with in the same way as standard Gettier cases. The reason for this is that there doesn't appear to be an abnormality in play in the causal explanation of why Barnes forms a true belief. True, there might have been such an abnormality—for instance, if Barnes had arrived at a true belief despite looking at a fake barn. But, of course, this is not to say that the abnormality does actually obtain.

Accordingly, Greco pursues a rather different strategy here, which relies essentially on his contextualist semantics for knowledge attributions. More specifically, Greco aims to show that in the context that obtained when we considered FAKE BARNS and in which we found it intuitive that Barnes does not know that he is looking at a barn, the semantic value of 'ability to form true beliefs about the presence of barns' is something like *ability to form true beliefs about the presence of barns whilst moving through an area featuring mostly indistinguishable fakes*. The idea here is that the cognitive ability's *E* include the wider area in which Barnes finds himself and in which fake barns predominate. Moreover, the ability's *C* are not restricted to circumstances in which Barnes happens to look out of the window just when he passes the one real barn, but also include circumstances in which he looks out of the window a little bit earlier or later and ends up looking at a fake barn. But now notice that, given that the *C* and *E* are thus understood, the majority of nearby worlds at which Barnes is in these *C* and *E* and forms a belief about the presence of a barn are worlds at which Barnes ends up looking at a fake barn. Since, at those worlds, Barnes ends up with a false belief that he is looking at a barn, it is not the case that Barnes attains true beliefs about the presence of barns at a high rate across nearby worlds in these *C* and *E*. On Greco's account of abilities, then, Barnes does not have the ability to form true beliefs about the presence of barns relative to these *C* and *E*. In consequence, in our context, the attribution of the relevant cognitive ability to Barnes is false and the denial of knowledge is true.

But now notice that, on Greco's account, sentences of the form 'S knows that *p*' and 'S does not know that *p*' do not express propositions unless (i) a practical reasoning context and (ii) what counts as actionable information in that context have been determined. While Greco effectively endorses this consequence of his view [Greco 2010: 79-80], it does appear to generate some trouble for him.³ To see this, notice that the description of FAKE BARNS appears to leave both of

these issues wide open. But given that this is so, the sentence 'Barnes does not know that he is looking at a barn' doesn't express a proposition in the context that obtained when you read *FAKE BARNS*. At the same time, you agreed (if only implicitly) with my assertion of 'I take it to be agreed that Barnes does not know that he is looking at a barn'. Now, this is rather puzzling. If Greco is right and 'Barnes does not know that he is looking at a barn' does not express a proposition in that context, we'd expect competent speakers of English like yourself to disagree with this assertion. After all, the object of agreement is a proposition. If 'Barnes does not know that he is looking at a barn' doesn't express a proposition, there simply is no suitable object for agreement. We'd expect competent speakers of English to notice this and to refuse to accept my assertion. If this isn't immediately obvious, compare: You are standing in a room with fifty paintings by a variety of artists. I walk in and assert: 'I take it to be agreed that this is a Picasso.' offering no indication of what 'this' is supposed to refer to. Given that you are a competent speaker of English, we would expect you to disagree with my assertion. If you decided to respond, we would not expect you to say 'Yes, you are right.' Rather, we'd expect you, for instance, to point out that before you can agree you will have to know which painting I have in mind. In other words, we'd expect you to spot the problem and refuse to accept the assertion. For that reason, the fact that you agreed with my assertion about Barnes suggests that 'Barnes doesn't know that he is looking at a barn' did express a proposition when I asserted 'I take it to be agreed that Barnes does not know that he is looking at a barn' in the context that obtained when you read *FAKE BARNS*. Greco thus faces a difficulty.

I assume that Greco's response to this difficulty is that there is a default practical reasoning context and a default standard for actionable information, which we are in play in contexts in which no explicit information about the practical reasoning context and what counts as actionable information are provided. If so, then it is the default practical reasoning context and standard for actionable information that are in play in the context that obtained when you agreed that Barnes does not know that he is facing a barn upon reading *FAKE BARNS*. For Greco's account of the absence of knowledge in *FAKE BARNS* to work, the *C* and *E* at issue in the default practical reasoning context and given default standards for actionable information must be wide enough to include the wider area Barnes finds

himself in as well as circumstances in which Barnes looks out of the window a bit earlier or later and ends up looking at a fake barn.

With these points in play, let's return to *FRANKFURT CLOCK*. Notice that, here too, we have not been given explicit information about the practical reasoning context and the standards for actionable information. Yet, again we have no difficulties in agreeing that Russell knows. We may thus assume that the default practical reasoning context and standard for actionable information obtain. Now, crucially, if, in *FAKE BARNS*, the *C* and *E* at issue in the default practical reasoning context and given default standards for actionable information are wide enough to include circumstances in which Barnes looks out of the window a bit earlier or later and ends up looking at a fake barn, then, in *FRANKFURT CLOCK*, the *C* and *E* that are in play by default will be wide enough to include circumstances in which Russell looks at the clock a bit earlier or later. Of course, in those cases, the demon shows his hand and Russell ends up with a false belief about the time. The difficulty for Greco is that that the majority of nearby worlds are worlds at which one of these circumstances obtains. It is thus not the case that Russell attains true beliefs about the time at a high rate across nearby worlds in the relevant *C* and *E*. On Greco's account of abilities, then, Russell does not have the ability to form true beliefs about the time relative to these *C* and *E*. The attribution of the relevant cognitive ability to Russell is false in the context that obtained when you read *FRANKFURT CLOCK* and, in consequence, the denial of knowledge is true in that context. Greco's account of *FAKE BARNS*, if successful, will lead him to pass the wrong verdict in *FRANKFURT CLOCK*. This is, of course, bad news for Greco.

3.2 Sosa

According to Sosa, abilities are dispositions to perform well. Roughly speaking, they are dispositions that issue performances that, in turn are likely to be successful. Let's see how Sosa unpacks this rough idea in a bit more detail.

Dispositions have three essential components: a constitution component (*CO*), a condition (or shape) component (*SH*) and a situation component (*SI*). Consider, for instance, a match's disposition to light when struck. Its constitution component comprises certain physical and chemical properties of the powdery head, its condition component includes, among other things, being dry, and its situation component includes being in oxygen [Sosa 2010: 465]. Dispositions also

correspond to trigger-manifestation conditionals. To see what that means, consider the disposition to light when struck once more. The trigger here is striking, the manifestation is lighting. Our match can be said to have this disposition, if and only if the conditional ‘if the match were struck, it would likely light’ is true of it, at least when the match satisfies *CO* and is in *SH* and *SI* [Sosa 2010: 466].

In case of the cognitive abilities at issue in *VE*, the performances issued are beliefs and the successes are truths [Sosa 2011: 1]. For instance, the ability to read the clock is a disposition that issues beliefs about the time. A given belief about the time is successful if true. Cognitive abilities also feature *CO*, *SH* and *SI* components. In case of the ability to read the clock, the *CO* are properties of the agent that constitute the seat of the disposition, the *SH* include, among other things, being awake and sober and the *SI* include there being enough light to take a reading and a normally functioning clock. Cognitive abilities also correspond to trigger-manifestation conditionals. For instance, one can be said to have the ability to read the clock if and only if the conditional ‘if one were to form a belief about the time, one’s belief would likely be true’ is true of one, at least when one satisfies the relevant *CO* and is in *SH* and *SI*.

Sosa analyses the because relation in terms of ability manifestation: a success is because of ability just in case the success manifests ability. Crucially, a success manifests a certain ability only if the ability’s *SH* and *SI* are satisfied [Sosa 2010: 470]. According to Sosa, when our agent’s ability to read the clock issues a true belief about the time in suitable *SH* and *SI*, his success manifests his clock-reading ability. It does so (in part) because the *SH* and *SI* are suitable. In contrast, when the *SH* or *SI* are not suitable the agent will not attain success because of ability, not even if the performance turns out to be successful.

With these points in play it is easy to see how Sosa can deal with standard Gettier cases: while the agent attains the relevant cognitive success, the relevant cognitive ability’s *SI* are not satisfied. For instance, consider once more the case of the agent who forms a true belief by taking a reading from a clock that stopped exactly twelve hours earlier. Since the clock is stopped, the *SI* of his clock-reading ability are not satisfied. In consequence, the agent’s true belief does not manifest his clock-reading ability. The agent falls short of attaining success because of ability. In this way, Sosa can explain the absence of knowledge in this case.

Just like Greco, Sosa agrees that Fake Barn cases cannot be dealt with in the same way as standard Gettier cases. In fact, Sosa acknowledges (i) that there are non-epistemic analogues of these cases in which (ii) the agents do attain success because of ability, and that (iii), for that reason, the agents in Fake Barn cases should be credited with success because of ability also. Unsurprisingly, then, Sosa thinks it would be a mistake to opt for the kind of account of these cases favoured by Greco and deny agents in Fake Barn cases even the possession of the relevant cognitive abilities [Sosa 2010: 469]. But given that Sosa accepts that agents in Fake Barn cases attain success because of ability, how can he account for the intuition that they lack knowledge?

In order to answer this question, Sosa introduces a distinction between animal knowledge and reflective knowledge. Animal knowledge is true belief that manifests a first-order ability, such as, for instance, the ability to spot barns. Reflective knowledge, in contrast, requires true belief that, in addition, manifests a second-order monitoring ability. Sosa's strategy is to argue that, while agents in Fake Barn cases have animal knowledge, they lack reflective knowledge. Crucially, according to Sosa, intuitions about knowledge track the presence and absence of *reflective* knowledge [Sosa 2010: 474 and Sosa 2011: 92-93]. As a result, if the strategy can be successfully implemented, the intuition of ignorance will be duly explained.

Let's take a closer look at how Sosa ventures to argue that agents in Fake Barn cases lack reflective knowledge. As a first step, he notes that, just like first-order abilities, second-order abilities, too, are dispositions that make success highly likely and are relative to shape and situational conditions. Crucially, according to Sosa, the *SI* of the second-order abilities include the following condition: one must be so situated that "sooner or later there would be tell-tale signs" if the first-order ability's *SH* or *SI* were unsuitable [Sosa 2010: 473]. The reason why agents in Fake Barn cases lack reflective knowledge is that it is precisely not the case that there would sooner or later be tell-tale signs if the *SI* of the relevant first-order ability were not satisfied. For instance, in *FAKE BARNs*, if the *SI* of Barnes's first-order barn-spotting ability were not satisfied, Barnes would be looking at an indistinguishable fake barn. In that case, however, there would be no indication of the fact that the *SI* for the first-order ability are not in place. So, even if Barnes's true belief that he is looking at a barn manifests first-order cognitive ability, it does not manifest second-order

ability. Barnes attains animal knowledge but falls short of reflective knowledge.

It is not hard to see that, if Sosa's story handles Fake Barn cases in a satisfactory manner, it is bound to get him into trouble with Frankfurt cases. Let's return to FRANKFURT CLOCK and ask what would be the case if the *SI* for the first-order ability were not satisfied. The answer is that Russell would have come down a minute earlier or later and the demon would have manipulated the clock to read '8:22' anyway. It is easy to see that there would not sooner or later be tell-tale signs for Russell to pick up on here. So, the *SI* for Russell's second-order ability are not satisfied. Russell falls short of reflective knowledge. Since intuitions about knowledge are said to track the presence and absence of reflective knowledge, Sosa's account incorrectly predicts an intuition of absence of knowledge here.

4 Conclusion

It comes to light that Frankfurt cases serve to generate a problem for some of the most promising accounts of knowledge on the epistemological market, to wit safety-based accounts and the two leading virtue epistemological accounts.

Does that mean that we should abandon these accounts and move on in our search for a workable account of knowledge? Perhaps. On the other hand, it may be that the problem here instantiates a deeper underlying difficulty, which affects a much wider range of (and perhaps all) accounts of knowledge that aim to analyse knowledge in terms of justified true belief and an anti-Gettier condition. I will not try to drive this point home here. Rather, I will rest content with stating the following challenge that any adequate such account of knowledge will have to meet, *viz.* to countenance an anti-Gettier condition that is strong enough to explain the absence of knowledge in Fake Barn cases, whilst, at the same time, being weak enough to allow the presence of knowledge in Frankfurt cases.

I'd like to add that meeting this challenge is by no means trivial. Most standard proposals in literature are prone to fail it in one way or another. Modal accounts, such as safety and sensitivity-based accounts, fail it because the conditions they impose are too strong: they predict absence of knowledge in Frankfurt cases. The same goes for no defeaters accounts.⁴ Other accounts fail in that the proposed conditions are too weak to explain the absence of knowledge in Fake

Barn cases. Some less sophisticated versions of virtue epistemology, i.e. versions that do not add a special story for Fake Barn cases, are a case in point, as are the closely related proper functionalist accounts. No false lemma and causal accounts have long been known to succumb to Fake Barn cases. If there exists a workable account of knowledge in terms of justified true belief and an anti-Gettier condition that gets both Fake Barn cases and Frankfurt cases right, epistemology still awaits its discovery.

Perhaps, rather than holding our breath for an account of knowledge that succeeds where its predecessors have failed, we should do agree with Pritchard and prepare to sacrifice at least one of the intuitions in Fake Barn cases and Frankfurt cases. Even so, we ought not dismiss either intuition simply out of hand. At the very least, we need an error theory for the intuition we sacrifice that is workable in the sense that it explains why intuition leads us astray in one of the two types of case, without incorrectly predicting the presence/absence of the intuition in the other. And it is worth noting that epistemology also still awaits its discovery.

So the conclusion that we might be inclined to draw at this stage is that vital work remains to be done in epistemology. But wait! There is another option. After all, it might be that, while there is an important difference between Fake Barn and Frankfurt cases, this difference cannot be explained without invoking the concept of knowledge. For instance, it may be that the way in which Russell and Barnes differ is just not be explicable without saying that Russell has knowledge whilst Barnes doesn't. One advantage of this interpretation of the cases is that we can avoid sacrificing intuitions, another that we do not need to hold our breath for further epistemological discoveries.

Of course, if there are genuine differences between these cases that cannot be explained without invoking the concept of knowledge, there is little hope for an analysis of knowledge in terms of justified true belief and an anti-Gettier condition. After all, any such analysis is then bound to make the wrong predictions either in Fake Barn cases or in Frankfurt cases. In that case, Fake Barn and Frankfurt cases would provide reason to favour a knowledge first approach to epistemology [Williamson 2000, 2010] over its traditional rivals.

Notes

¹ It may be worth noting that, in more recent work, Pritchard has moved to a full-blown sliding scale account of the tolerance for error at nearby worlds. The

idea is that at nearby worlds closest to actuality, safety tolerates no false beliefs and the further we move away from actuality the more false beliefs safety is compatible with. However, Pritchard has not given a precise statement of safety in recent work. Note also that, in recent references to “versions of the safety principle” in the literature, Pritchard [2012b: 170, n.3, 2013: n.10] lists the above as the last version in his own work.

² [Author 2009a]. I owe inspiration for this case to Frankfurt’s famous [1969] cases. For alternative counterexamples to the safety condition see [Comesaña 2005, Neta and Rohrbaugh 2004]. In [Author 2009a], I also argue that my Frankfurt case makes for a better counterexample to the safety condition than these alternatives.

³ In fact, there is reason to believe that it is a vital feature of our concept of knowledge that it is applicable without information about the specificities of some practical reasoning context [Craig 1990].

⁴ If this isn’t immediately obvious, note that both of the following constitute defeaters for Russell’s belief that it is 8:22: (i) ‘The grandfather clock is under the control of an agent who wants you to believe that it is 8:22 when you look at it.’ (ii) ‘The grandfather clock would have displayed ‘8:22’ even if it hadn’t been 8:22.’

References

Author 2009a. Author’s work.

Author 2009b. Author’s work.

Comesaña, J. 2005. Unsafe knowledge. *Synthese*, 146:395–404.

Craig, E. 1990. *Knowledge and the State of Nature. An Essay in Conceptual Synthesis*. Clarendon Press, Oxford.

Frankfurt, H. 1969. Alternate possibilities and moral responsibility. *The Journal of Philosophy*, 66:829–39.

Gettier, E. 1963. Is justified true belief knowledge? *Analysis*, 23:121–23.

Goldman, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy*, 73:771–91.

Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press, Cambridge.

Greco, J. 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research*, 85:1–26.

Neta, R. and Rohrbaugh, G. 2004. Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, 85:396–406.

- Nozick, R. 1981. *Philosophical Explanations*. Oxford University Press, Oxford.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press, Oxford.
- Pritchard, D. 2007. Anti-luck epistemology. *Synthese*, 158:277–297.
- Pritchard, D. 2008. Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, 58:437–47.
- Pritchard, D. 2009. Safety-based epistemology: Whither now? *Journal of Philosophical Research*, 34:33–45.
- Pritchard, D. 2012a. Anti-luck virtue epistemology. *The Journal of Philosophy*, 109:247–79.
- Pritchard, D. 2012b. In defence of modest anti-luck epistemology. In Becker, K. and Black, T., editors, *The Sensitivity Principle in Epistemology*. Cambridge University Press, Cambridge.
- Pritchard, D. 2013. Knowledge cannot be lucky. In Steup, M., Turri, J., and Sosa, E., editors, *Contemporary Debates in Epistemology*. Blackwell, Oxford, 2 edition.
- Pritchard, D., Millar, A., and Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford University Press, Oxford.
- Russell, B. 1948. *Human Knowledge: Its Scope and Limits*. Allen & Unwin, London.
- Sosa, E. 1999. How to defeat opposition to moore. In Tomberlin, J., editor, *Philosophical Perspectives*. Blackwell, Oxford.
- Sosa, E. 2010. How competence matters in epistemology. *Philosophical Perspectives*, 24:465–475.
- Sosa, E. 2011. *Knowing Full-Well*. Princeton University Press, Princeton/NJ.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press, Oxford.
- Williamson, T. 2010. Knowledge first epistemology. In Bernecker, S. and Pritchard, D., editors, *The Routledge Companion to Epistemology*, pages 208–18. Routledge, New York.