CrossMark

Episteme, 11, 4 (2014) 411–422 © Cambridge University Press
doi:10.1017/epi.2014.21

UNREFLECTIVE EPISTEMOLOGY

CHRISTOPH KELP
christoph.kelp@hiw.kuleuven.be

ABSTRACT

Virtue epistemological accounts of knowledge claim that knowledge is a species of a broader normative category, to wit of success from ability. Fake Barn cases pose a difficult problem for such accounts. In structurally analogous but non-epistemic cases, the agents attain the relevant success from ability. If knowledge is just another form of success from ability, the pressure is on to treat Fake Barn cases as cases of knowledge. The challenge virtue epistemology faces is to explain the intuitive lack of knowledge in Fake Barn cases, whilst holding on to the core claim that knowledge is success from ability. Ernest Sosa's version of virtue epistemology promises to rise to this challenge. Sosa distinguishes two types of knowledge, animal knowledge and reflective knowledge. He argues that while animal knowledge is present in Fake Barn cases, reflective knowledge is absent and ventures to explain the intuition of ignorance by the absence of reflective knowledge. This paper argues that Sosa's treatment of Fake Barn cases fails as it commits Sosa to a number of highly counterintuitive results elsewhere in epistemology.

1. INTRODUCTION

What is knowledge? It will come as no surprise that this is one of the core questions in the theory of knowledge. What is perhaps more surprising is how difficult it is to provide a satisfactory answer to it. Since Gettier's celebrated (1963) counterexamples to the then received view that knowledge is justified true belief, the epistemological literature has featured a long list of proposed accounts of knowledge and an almost equally long list of Gettier-style counterexamples to these accounts.

That said, some promising candidates have surfaced in recent years. Among them are virtue epistemological accounts. The core idea of virtue epistemology (VE) is that knowledge is a species of a broader normative genus, to wit, success from ability or competence. More specifically, virtue epistemologists maintain that knowledge is cognitive success (true belief) that is attained because of cognitive ability. This account of knowledge, they argue, will enable us to solve a number of core problems in epistemology, including the Gettier problem.[1]

VE's core idea carries some promise. To see this, consider first Chisholm's (1966) classic case:

1 Another central problem VE has been claimed to solve is the value problem. For more on this, see e.g. Greco (2010a, 2010b) and Riggs (2002).

> *Sheep*. Mutton, a reliable sheep-spotter, sees a sheep-shaped object and comes to believe that there is a sheep ahead. Unbeknownst to Mutton, the object he is looking at is not a sheep but a dog in fancy dress. At the same time, Mutton's belief is nonetheless true as there happens to be a sheep hidden from view behind the dog.

Mutton's belief falls short of knowledge. According to a widely accepted diagnosis, the reason for this is that it is pure luck that his belief is true. Notice also that Mutton's cognitive performance is tarnished as it involves him mistaking a non-sheep for a sheep. We may thus want to say that Mutton is lucky to attain cognitive success in the sense that he succeeds despite his tarnished cognitive performance.

Crucially, second, it is independently plausible that success because of ability contrasts with lucky success of this kind, i.e. success despite tarnished performance. By way of evidence, consider Archer who produces a competent shot which is first blown off target by a freak gust of wind and then brought back on target by a second freak gust. Even though Archer's shot is successful, it is not successful because of ability. Notice also just how appealing the parallel diagnosis is: the reason why Archer is not successful because of ability is that it is pure luck that his shot hits the target. Moreover, Archer's athletic performance is tarnished also as it involves a shot that has been blown off target. Archer, too, is lucky to attain the relevant success in the sense that he succeeds despite his tarnished performance.

If, as VE would have it, (VE1) knowledge is success from ability; if (VE2) success from ability is incompatible with success despite tarnished performance; and if (VE3) Gettier cases are cases of success despite tarnished performance, it looks as though VE has what it takes to explain why agents in Gettier cases lack knowledge.

But now consider another prominent Gettier case, which is due to Alvin Goldman (1976) (who in turn attributes it to Carl Ginet):

> *Barns*. Barnes, a reliable barn-spotter, drives through the countryside, sees a barn in the field to the right and forms a true belief that he is facing a barn. Unbeknownst to Barnes, he is looking at the only real barn in a field otherwise full of fake barns that are so cleverly constructed as to be indistinguishable from real barns from Barnes's position on the road.

Barnes's belief falls short of knowledge. Moreover, the above diagnosis is still plausible: it is pure luck that Barnes believes truly here. In these respects, *Barns* is analogous to *Sheep*.

At the same time, there is an important structural difference between *Barns* and standard Gettier cases like *Sheep*: in *Barns* the agent's performance is *not* tarnished.[2] Unlike Mutton, who mistakes a non-sheep for a sheep, Barnes does not mistake a fake barn for a real barn. Of course, he might so easily have looked at a fake instead in which case he would have made a similar mistake as Mutton. However, as a matter of fact, he does not. What is going on in *Barns*, then, is that Barnes's cognitive performance is untarnished but *might so easily have been tarnished*. As a result, (VE3) is false: Gettier cases need not be cases of success despite tarnished performance. Rather cases in which the agent's successful performance is untarnished but might easily have been tarnished will do as well.

This means trouble for VE. Here is why. While agents who succeed despite tarnished performance *generally don't* succeed because of ability, there is excellent evidence that

2 For alternative accounts of this difference see Hetherington (1999) and Pritchard (2008).

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

agents who succeed via untarnished performances that might easily have been tarnished *generally do* succeed because of ability. Just think of a case in which Archer produces a successful shot at a normal target when shooting at the only unsabotaged target at a shooting range otherwise full of sabotaged targets; think of Chef who successfully prepares a tasty omelette but happened to take the only salt shaker in which the salt wasn't replaced by sugar; or think of Artist who successfully produces a beautiful monochrome after having taken the only can in which the colour wasn't replaced by acid. In all of these cases, the agents produce performances that are untarnished but might easily have been tarnished. At the same time, it is intuitively clear that they succeed because of ability. If, as VE would have it, knowledge is a species of success because of ability, we have excellent reason to believe that VE will predict, incorrectly, that Barnes knows that he is facing a barn. If so, VE succumbs to Gettier-style counterexample after all.[3]

Can VE avoid this difficulty? That is to say, is it possible for virtue epistemologists to treat cases like *Barns* analogously to the cases of Archer, Chef and Artist as cases of success because of ability, whilst explaining the intuition that Barnes lacks knowledge? One might think that the answer to these questions has to be 'no'. However, Ernest Sosa has recently offered a version of VE that suggests otherwise. His account will be the main focus of this paper. I will argue that this answer, while initially promising, does not stand up to closer scrutiny. In order to achieve this, it will be necessary to first outline Sosa's virtue epistemology and his account of *Barns* (and similar cases) in more detail.

2. SOSA'S VIRTUE EPISTEMOLOGY

2.1 ... *In outline*

Sosa (2007, 2010, 2011) offers an account of the normativity of performances with an aim. According to Sosa, performances with an aim fall under 'AAA structure'. This means they can be evaluated in terms of whether they attain the aim (i.e. whether they are successful), whether they are adroit (i.e. whether they manifest competence) and whether they are apt (i.e. whether they are successful because competent). For instance, the aim of a shot in target archery is to hit the target. To say that a shot in target archery

3 Pritchard et al. (2010), Pritchard (2012). Some champions of VE (most notably Greco 2007, 2010a) respond by holding that Barnes does not even have the ability to form true beliefs about the presence of barns in fake barn county. However, this move is not particularly attractive. Among the undesirable consequences it brings in its trail are: (i) It fails to treat like cases alike. (ii) It threatens to render Mutton epistemically better off than Barnes. After all, Mutton but not Barnes has and exercises the kind of cognitive ability required for knowledge. Relatedly, (iii) Greco's move threatens to render Barnes's belief unjustified. According to a plausible virtue epistemological account of justified belief, a belief is justified just in case it is the product of the kind of cognitive ability required for knowing. Given that, in *Barns*, Barnes does not have the relevant ability, his belief turns out to be unjustified. Finally, (iv) if it is not entirely *ad hoc*, it must be supported by an independent argument. However, this threatens to smuggle in covert additional conditions on knowledge that do part of the anti-Gettier duty. If so, VE is compromised, at least in the sense that it cannot be the full epistemological story about knowledge. This might in turn mean that other parts of VE's story will no longer work in the way envisaged. In particular, I am thinking here of the kind of account of the value of knowledge favoured by Greco (e.g. 2010a, 2010b), which crucially depends on VE's condition being not only necessary but also sufficient for knowledge.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

falls under AAA structure is to say that it can be evaluated in terms of whether it hits the target, whether it is competent and whether it hits the target because it is competent.

Competences, according to Sosa, are dispositions to perform well. That is to say, competences are dispositions that issue performances that, in turn, are likely to be successful.

Dispositions have three essential components: a *constitution component* (CO), a *condition* (*or shape*) *component* (SH) and a *situation component* (SI). For instance, a match's disposition of being flammable has as its constitution component certain physical and chemical properties of the powdery head. These properties constitute the seat of the disposition. 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. In the case of our match's flammability, this conditional is: if it were struck it would likely light. The idea here seems to be that a disposition corresponds to a trigger-manifestation conditional if and only if that conditional must be true of the object satisfying CO and being in SH and SI. To say that the match's flammability corresponds to the conditional 'if it were struck it would likely light' is to say that this conditional must be true of any match the powdery head of which has the relevant physical and chemical properties, which is dry, etc., and which is in oxygen, etc. (Sosa 2010: 466).

Since competences are dispositions, they also feature CO, SH and SI components and correspond to trigger-manifestation conditionals of the form: if the disposition were to issue a performance, the performance would likely be successful. In the case of archery competence, the CO are certain properties of the agent, including more basic motor and cognitive competences, the SH include being awake and sober and the SI include there being enough light and normal winds. Furthermore, since the performances the disposition issues here are shots and successes hits, it corresponds to the following trigger-manifestation conditional: if the agent were to take a shot, he would likely hit the target. So, a competent archer is an agent who has certain properties such that when awake, sober, etc. and there is enough light, winds are normal, etc., he would likely hit the target were he to shoot at it (Sosa 2010: 466).

Sosa also distinguishes between three different levels of competence: first, the *constitutional competence* which consists of the constitutional component of the agent's disposition (CO); second, the *inner competence* which combines the constitutional component with the shape component (CO and SH); and, third, the *complete competence* which, besides the constitutional and the shape components, also includes the situational component (CO, SH and SI). On this picture, it is of course the complete competence that corresponds to the trigger-manifestation conditional (Sosa 2010: 465).

According to Sosa, a performance is successful because competent or apt if and only if its success manifests the agent's competence. With the above account of competence in play, we can say a bit more about what that means. Crucially, aptness is unpacked in terms of the notion of a *complete* competence: 'A performance is apt', Sosa tells us, 'if and only if its success manifests a complete competence' (Sosa 2010: 470). Aptness thus requires that the relevant disposition's CO, SH and SI are in place. Suppose, for instance, that an agent is a competent archer in virtue of having certain properties such that when awake, sober, etc. and there is enough light, winds are normal, etc., he would likely hit the target were he to shoot at it. In that case, any shot he fires will be apt only if he is in fact awake, sober, etc. and there is in fact enough light, winds are normal, etc. Any shot he fires while drunk, in darkness or in abnormal winds will not qualify as apt.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

Crucially, according to Sosa, beliefs are performances with an aim. More specifically, the aim of belief is truth (Sosa 2011: 1). In consequence, beliefs too fall under AAA structure. They can be evaluated in terms of whether they attain the aim of truth, whether they are adroit or competent, and whether they are apt or true because competent. Unsurprisingly, Sosa's suggestion is that knowledge is apt belief.

This account of competences can be applied to the kinds of competences required for knowing ('cognitive competences') in the following way: Cognitive competences are dispositions of agents featuring CO, SH and SI. Since the performances these dispositions issue are beliefs and successes are truths, they correspond to the following kind of trigger-manifestation conditional: if the agent were to believe a proposition within a certain range, he would likely believe truly. For instance, an agent has a cognitive competence to spot barns just in case he has certain properties such that when in suitable SH (awake, sober, etc.) and SI (sufficient light, etc.), if he were to believe that a barn is present, his belief would likely be true. Of course, in case of cognitive competence, too, the three levels of competence (constitutional, inner and complete) can be distinguished. Crucially, for cognitive competences, too, aptness requires that the agent's success manifests a complete competence. In consequence, any belief formed in untoward SH or SI is bound to fall short of knowledge, no matter whether it is adroit and/or turns out to be true.

2.2 *... And Fake Barn cases*

With Sosa's virtue epistemology in play, let's return to *Barns*. How can Sosa accommodate the intuition of ignorance in this case? Sosa tackles this question by introducing a distinction between *animal knowledge* and *reflective knowledge*. While animal knowledge is first-order apt belief, reflective knowledge is meta-apt belief. Meta-apt belief involves the manifestation of a second-order cognitive competence, which is, in essence, a monitoring competence. More specifically, according to Sosa, the relevant second-order competences are 'default competences' in the sense that they assume (by default) that the complete first-order competence is present, unless there are indications to the contrary (Sosa 2010: 473). Like their first-order cousins, second-order competences come in three varieties: constitutional, inner and complete. Here, too, the constitutional competence consists in the seat of the competence (CO) only, the inner combines the seat and the inner component (CO and SH) and the complete competence includes, in addition, the situational component (CO, SH and SI). According to Sosa, the SI of second-order cognitive competences include the following condition:

> *Tell-Tale Signs* (*TTS*). '[O]ne must be so situated that sooner or later there would be tell-tale signs if the relevant [complete] first-order competence were absent' (Sosa 2010: 473).

For a belief to be meta-apt it must manifest a *complete* second-order competence. A belief will thus be meta-apt only if TTS is satisfied. It is easy to see that, in consequence, a belief can be apt without being meta-apt. This will happen when the agent rises to the level of complete first-order competence but is not so situated that there would be tell-tale signs if he didn't. In that case the agent will acquire animal knowledge, without acquiring reflective knowledge.

According to Sosa, this is exactly what happens in *Barns*: Sosa accepts that Barnes rises to the level of complete competence at the first order and thus acquires animal knowledge.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

On the second order, however, things look differently: if Barnes's first-order competence were absent here, this would be because Barnes is looking at a fake barn. In that case, there would be nothing about the situation that would alert Barnes to this fact. That means that Barnes does not satisfy TTS. As a result, Barnes does not rise to the level of complete competence at the second order. Barnes's belief does not qualify as reflective knowledge (Sosa 2010: 473–4).

Crucially, according to Sosa, intuitions about knowledge track the presence and absence of *reflective* knowledge. Since Barnes does not attain reflective knowledge, the intuition of ignorance in *Barns* is duly explained.

Before moving on, I would like to discuss one immediate problem that arises for Sosa's account at this stage. To bring it out, consider first the following case:

> *Barns**. Barnes is driving through a part of the country featuring mostly fake barns. In order to prevent people from forming false beliefs, the authorities have seen to it that the fake barns (and only the fake barns) are painted blue. In addition, they have placed signs at both ends of Fake Barn County alerting drivers to the fact that all and only blue structures that look like barns are in fact mere façades. Unfortunately, the sign at the end of the road at which Barnes enters Fake Barn County had been vandalised the night before. As a result, whilst driving through Fake Barn County, Barnes does not know that the blue structures are fake barns or that most structures that look like barns are in fact mere façades. Upon seeing one of the few (non-blue) real barns, Barnes forms the belief that he is facing a barn. When Barnes exits Fake Barn County he notices the sign.

Consider Barnes's belief about the presence of a barn at the time at which he first forms it (t_0). Intuitively, at t_0, Barnes's belief falls short of knowledge. Sosa wants to explain this intuition in terms of the absence of reflective knowledge and that, in turn, in terms of TTS's not being satisfied. Recall that, according to TTS, the SI of a second-order competence will be satisfied only if the agent is so situated that were the first-order competence to be absent, sooner or later there would be tell-tale signs alerting him to this fact. Crucially, there is reason to believe that TTS is satisfied in *Barns** at t_0. After all, if, at t_0, the first-order competence were absent, this would be because the structure Barnes is looking at is a blue fake barn – or so we may assume. In that case, however, the relevant tell-tale signs do present themselves sooner or later, *viz.* upon exiting Fake Barn County when Barnes passes by a sign telling him that the blue barn-shaped objects are mere façades. In *Barns**, then, Barnes is so situated that if the first-order competence were absent, sooner or later there would be tell-tale signs. TTS turns out to be satisfied. As a result, Sosa's envisaged explanation of the absence of reflective knowledge at t_0 in terms of TTS fails.

Fortunately, this problem can be solved. Sosa can simply strengthen TTS in the following way. Rather than allowing the tell-tale signs to present themselves *sooner or later*, Sosa can require that the relevant tell-tale signs present themselves *soon enough*. If, in addition, it can be argued that in *Barns**, the tell-tale signs would not present themselves soon enough, Sosa's explanation of Barnes's lack of knowledge will be reinstated.

Of course, the question that immediately arises at this stage is what counts as soon enough. Curiously, there is reason to believe that Sosa cannot allow the extension of 'soon enough' to extend to any times beyond the time of belief-formation. To see this notice that we can easily set up *Barns** in such a way that Barnes notices the sign at t_1, i.e. immediately after forming his belief about the presence of the barn at t_0. In that

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

case, still Barnes's belief intuitively falls short of knowledge at t_0. However, if t_1 falls within the extension of 'soon enough', then, at t_0, Barnes will satisfy even the strengthened version of TTS, for reasons already discussed. Once again, the envisaged explanation of the absence of reflective knowledge at t_0 in terms of TTS fails.

I take these considerations to indicate that, in order to successfully explain the absence of reflective knowledge in cases like *Barns*, TTS will have to be restricted to the time of belief-formation. For that reason, in the remainder of the paper I will assume that, properly understood, Sosa's crucial requirement is the following:

> TTS*. One must be so situated that, at the time of belief formation, there would be tell-tale signs if the relevant complete first-order competence were absent.

3. PROBLEMS FOR SOSA

3.1 Frankfurt cases

The first problem for Sosa I'd like to discuss concerns the second-order competences required by reflective knowledge. Sosa's account of these competences as competences that respond to the relevant tell-tale signs seems plausible enough. What is less clear to me, however, is that such competences must be manifested in order to acquire knowledge, at least if the SI include TTS*. To see this consider first the following case:

> *Frame*. Killer plans on murdering Victim. Controller, an agent of immense power, wants Killer to be convicted for the crime. In order to ensure this he is prepared to kill Victim himself and convincingly frame the deed on Killer if Killer decides not to carry out his plan. Similarly, he will plant incriminating evidence against Killer if Killer does not leave enough clues. Killer goes ahead and murders Victim, leaving enough clues for Detective to mount a watertight case against Killer. Controller never shows his hand.[4]

When Detective has completed his case against Killer, intuitively, he knows that Killer did the deed. At the same time, Detective's belief does not qualify as reflective knowledge. If Detective's complete first-order competence had been absent, this would have been because Killer decided not to carry out his plan or didn't leave enough clues. In that case, however, Controller would have stepped in and murdered Victim and/or planted the evidence incriminating Killer. In that case there would not have been tell-tale signs that Detective's first-order competence is absent. Given that absence of first-order competence would not have been accompanied by tell-tale signs, Detective's second-order competence does not satisfy TTS*. By Sosa's lights, his belief does not qualify as reflective knowledge.[5] This is, of course, bad news for Sosa. After all, Detective's belief intuitively

4 It should be said that I owe inspiration for this example to Harry Frankfurt (1969) who uses examples with a similar structure to argue against the view that moral responsibility requires the ability to do otherwise. For an argument that this kind of case poses a problem for safety-based accounts of knowledge see Kelp (2009).

5 Notice that this will follow even on TTS. After all, we may assume that, had Controller showed his hand, he would have seen to it that the relevant tell-tale signs would not surface sooner or later either.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

qualifies as knowledge. If intuitions about knowledge track the presence and absence of reflective knowledge, it looks as though we have a counterexample to Sosa's account.[6]

As I see it the problem here is that the second-order competences are construed in too strong a manner. More specifically, the difficulty arises from TTS*, the requirement that one be so situated that there would be tell-tale signs if one's first-order competence were absent. What *Frame* highlights is that one can have knowledge even when absence of competence would not have been accompanied by tell-tale signs.

One promising way of avoiding this problem is by weakening TTS*. Instead of requiring that tell-tale signs would be present if the first-order competence were absent, we might think of the second-order competences as competences that are such that if they were to *detect* tell-tale signs *as such*, this would duly be taken into account (henceforth 'TTS**'). Notice that, on TTS**, Detective has the complete second-order competence. After all, if Detective were to detect anything fishy, he would suspend judgment. In that case, Detective does acquire reflective knowledge after all. The counterexample can be avoided.

Unfortunately, this alternative comes at a cost for Sosa. There is now little hope for his intended explanation of the intuition of ignorance in *Barns*. After all, here too, TTS** is satisfied: if Barnes were to detect tell-tale signs as such, they would be duly taken into account. For instance, were Barnes to look at a broken façade that exposes the structural elements behind it, he would believe that he is not looking at a real barn. In consequence, if TTS* is replaced by the weaker TTS**, Sosa's explanation of the intuition of ignorance in *Barns* will no longer work.

3.2 *Lottery cases*

Let's look at yet another case:

> *Lottery*. Gambles tells his close friend Raffles that he has recently bought a ticket in the state lottery, *L*. Based on the fact that the odds against winning are so very high, Raffles comes to truly believe that Gambles's ticket won't win *L*.

Although the probabilistic evidence supporting Raffles's belief is excellent and the belief is true, Raffles does not know that his friend's ticket won't win *L*.

To the best of my knowledge, Sosa does not discuss cases like *Lottery* in his published work.[7] That is not to say, however, that they don't pose a problem for him. On the contrary, there is excellent reason to think that Sosa will struggle with these cases. To see this, let's first ask whether Raffles's cognitive success manifests a complete first-order competence here. Before answering this question, let's ask whether Raffles rises to the level of complete competence in this case. Very plausibly, the answer to this question is 'yes'.

6 Note that we can set the case up in such a way that the closest worlds at which Controller shows his hands are far-off. This will make things even trickier for Sosa. After all, in that case, any temptation to regard manifestation of the complete second-order competence as necessary for knowledge is further weakened.

7 Note that Greco (2003: §I.B) contains a brief discussion of Sosa on lottery cases. In fact, Greco also notes that the intuition of ignorance cannot be explained in terms of absence of reflective knowledge here. That said, Greco's argument focuses mainly on an account of such cases in terms of a safety condition on knowledge, which Sosa was sympathetic to at the time, and does not provide a detailed argument that an explanation in terms of reflective knowledge will not do.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

After all, Raffles's belief is produced by disposition that in his SH (Raffles is awake, alert, etc.) and SI (the lottery is fair, etc.) corresponds to the trigger-manifestation conditional: if Raffles were to believe that Gambles's ticket will lose, his belief would very likely be true. So the disposition that produces Raffles's belief qualifies as a complete cognitive competence. If so, his cognitive performance (the belief issued by this disposition) manifests a complete cognitive competence in this case. Notice, furthermore, that his cognitive performance is also successful. After all, his belief that Gambles's ticket won't win the lottery is true.

With these points in play, let's return to the question whether Raffles's cognitive success manifests this complete competence in this case. Answering it is actually a bit tricky. The reason for this is that it does not follow from Sosa's account of aptness that cases in which performances are both successful and manifest a complete competence are cases of apt performances, i.e. cases of performances in which the *success* manifests a complete competence. That said, it is quite hard to see how cases of successful performance that manifest a complete competence could still fail to be cases in which the performance is apt. Think, for instance, of Archer once more. Consider a case in which he is taking a shot in perfect conditions (peak concentration, no winds, etc.) that hits the target. What reason could there possibly be to deny that his shot is apt? I, for one, can't think of any. If so, it is fair to assume that cases of performances that are successful and manifest complete competences are also cases of apt performances. Given Raffles's belief that Gambles's ticket won't win the lottery is both successful (true) and manifests a complete cognitive competence, there is thus reason to believe that it is also apt. In other words, there is reason to believe that it qualifies as animal knowledge.

Does Raffles's true belief also qualify as reflective knowledge? Does it manifest a complete second-order competence? First, note that we may assume that Raffles's belief does manifest the second-order monitoring competence of the sort required for reflective knowledge. By way of evidence, notice that we may assume that Raffles is such that had tell-tale signs been detected by the competence, they would have duly been taken into account. The crucial question, then, is whether Raffles's rises to the level of *complete* competence at the second-order as well. This question is also tricky to answer. One reason for this is that we need to know whether TTS* is satisfied. And, in order to ascertain this, we need to know what would have been the case if Raffles's complete first-order competence had been absent.

One might think that Raffles's complete first-order competence would have been absent in a case in which Gambles's ticket wins *L*. Since, in that case, there would not have been the required tell-tale signs, TTS* is not satisfied. Raffles's belief does not manifest a complete second-order competence and thus falls short of reflective knowledge. Unfortunately, there is reason to think that this doesn't work: Raffles's complete first-order competence would *not* have been absent in a case in which Gambles's ticket wins *L*. After all, even in that case, Raffles still satisfies the trigger-manifestation conditional: if he were to believe that Gambles's ticket loses, his belief is very likely true. For that reason, the SI of Raffles's first-order competence should include the case in which Gambles's ticket wins. It's just that, in this case, the unlikely possibility of failure materialises.

If the case in which Raffles's first-order competence would have been absent is not one in which Gambles's ticket wins *L*, presumably it is either a case (i) in which Raffles has no idea that Gambles owns a ticket to begin with or else a case (ii) in which Raffles knows that Gambles does not own a ticket. In both cases, there would be tell-tale signs that he

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

does not have the complete first-order competence (specifically: that he is in SI such that the relevant trigger-manifestation conditional is not true of him). In case (i) the tell-tale sign is the fact that Raffles has no idea that Gambles owns a ticket, in case (ii) it's the fact that Raffles knows that Gambles doesn't own a ticket. There is thus reason to believe that TTS* is satisfied.

This leads us immediately to another difficulty in ascertaining whether Raffles's second-order competence is complete here. Completeness requires that all SH and SI of the second-order competence are satisfied. So far, we have only shown that one such SI, TTS*, is. What remains to be shown is that the other SH and SI are satisfied as well. This task is complicated by the fact that Sosa typically leaves the list of SH and SI open-ended. So how can we possibly hope to answer this question?

The answer is: 'By investigating trigger-manifestation conditionals.' If relative to his SH and SI, Raffles satisfies the trigger-manifestation conditional that corresponds to the second-order competence, Raffles rises to the level of complete competence. What then is the trigger-manifestation conditional for Raffles's second-order competences? To answer this question, recall that Sosa construes the second-order competences as default competences that assume (by default) that the complete first-order competence is present. Given that this is so, the trigger-manifestation conditional corresponding to such second-order competences will be conditionals of the form: if – in suitable SH and SI – the second-order competence were to assume that the complete first-order competence is present, then it would likely be present. Crucially, there is excellent reason to believe that, in *Lottery*, the relevant trigger-manifestation conditional is true of Raffles. After all, we may assume that Raffles is in maximally favourable SH: he is awake, focused, at peak concentration, etc. Moreover, as we have already seen, TTS* is satisfied: Raffles is so situated that if the complete first-order competence were absent, then there would be tell-tale signs. Now consider: if – whilst Raffles is awake, focused, at peak concentration, etc. and so situated that if the complete first-order competence were absent, there would be tell-tale signs – Raffles's second-order competence were to assume that the complete first-order competence is present, it would likely be present. This trigger-manifestation conditional is very plausibly true. Since relative to Raffles's SH and SI, the trigger-manifestation conditional holds, Raffles's second-order competence is complete. In consequence, Raffles attains not only animal knowledge, but also reflective knowledge. Sosa's account makes the wrong predictions in cases like *Lottery*, too.

3.3 Dreaming scepticism

My last worry with Sosa's account arises from another case which Sosa himself discusses. It takes the following shape:

> *Dream.* Late at night, Wakes is sitting at home on his sofa and truly believes that he is awake. However, Wakes might very easily have gone to bed earlier in which case he would now lying in his bed dreaming that he is at the ball game, travelling Asia or sitting at home on his sofa. In that case, Wakes would also have believed, albeit now falsely, that he is awake. (Sosa 2010: 467)

Sosa grants that, intuitively, Wakes knows that he is awake. Given that Sosa is right in that intuitions about knowledge track the presence and absence of reflective knowledge, Sosa

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

had better establish that Wakes's belief qualifies as *reflective* knowledge. Now, the question is whether Sosa can secure this result. This task is complicated by the fact that Wakes is in a similar situation to Barnes in that both of them might very easily have acquired false beliefs instead without being in a position to know this. The challenge Sosa thus faces is to show that Wakes's belief qualifies as reflective knowledge without compromising his earlier account of Barnes's ignorance in terms of a lack of reflective knowledge.

Sosa takes on the challenge in the following passage:

> According to [Austin and Descartes], our dreams are incoherent, or have a dreamlike quality (an Austinian quality that may or may not be something other than the Cartesian incoherence). It's just that once asleep we are no longer in proper shape to take that into account. So, once asleep, we lack the relevant competence. Even if, in line with Austin and Descartes, we retain the situational element of that complete second-order competence, we lose an essential inner component in any case. Once asleep, we are in bad shape epistemically, with a consequent loss of competence. But here's the important point: *[Wakes's] incompetence when asleep is of no relevance to whether he has the complete second-order competence when awake and alert*. So, this is how [Wakes] is more fortunate than [Barnes]. This is why he attains a level of knowledge denied to [Barnes]. (Sosa 2010: 474)

The question remains how good Sosa's story is. In order to answer this question, consider the following case:

> *Chocolates*. Sweets is sitting in front of a box labelled 'chocolates'. He has just picked up one of them and is about to eat it: the chocolate is already touching his lips.[8] Unbeknownst to Sweets, the chocolate he is about to eat is the only chocolate in a box full of items that look just like chocolates but aren't.

Let's ask whether Sweets knows that he is about to eat a chocolate. Very plausibly the answer to this question is 'no'. The case is structurally analogous to *Barns*. Notice also that, as a result, Sosa had better analyse the case as one in which Sweets has animal knowledge but lacks reflective knowledge.

Of course, so far there is no special problem for Sosa. But now suppose we add:

> *Chocolates, Continued*. The fake chocolates are in fact drugs that upon contact with the mucous membrane would instantly put Sweets to sleep and would cause him to dream that he is sitting in front of a box of chocolates and is in the process of eating one of them.

To begin with, notice that this addition does not alter the intuition that Sweets does not know that he is about to eat a chocolate. With this point in play, here, then, is the crucial question: Does Sweets have his complete second-order competence? Given that dreams provide the required tell-tale signs, it would seem that Sweets is fine on this front. Moreover, given that inability to recognise dreams as such due to diminished competence leaves the second-order competences intact, there are no problems here either. In this respect, *Chocolates* is just like *Dream*. If so, and if, in *Dream*, Wakes retains his

8 We may add that they haven't come in contact with his taste buds yet and so Sweets does not yet have a gustatory experience as of a chocolate.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press

second-order competence and acquires reflective knowledge through it, then the same holds for *Chocolates*.

4. CONCLUSION

Sosa's version of VE promises to explain the intuition that agents in Fake Barn cases lack knowledge whilst granting that they attain successes because of ability. This paper has shown that Sosa's account of Fake Barn cases, whilst initially attractive, is ultimately unsuccessful as it comes at too high a cost: it leads him to mistaken predictions in Frankfurt-style as well as lottery cases, and, what's even worse, forces him down the road of dreaming scepticism. For that reason, I conclude, Sosa joins the ranks of those who have failed to provide an adequate solution to the Gettier problem.[9]

REFERENCES

Chisholm, R. 1966. *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
Frankfurt, H. 1969. 'Alternate Possibilities and Moral Responsibility.' *Journal of Philosophy*, 66: 829–39.
Gettier, E. 1963. 'Is Justified True Belief Knowledge?' *Analysis*, 23: 121–3.
Goldman, A. 1976. 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy*, 73: 771–91.
Greco, J. 2003. 'Knowledge as Credit for True Belief.' In M. DePaul and L. Zagzebski (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 111–34. Oxford: Oxford University Press.
——. 2007. 'The Purpose of Knowledge and the Nature of Ability.' *Philosophical Issues*, 17: 57–69.
——. 2010a. *Achieving Knowledge*. Cambridge: Cambridge University Press.
——. 2010b. 'The Value Problem.' In D. Pritchard and S. Bernecker (eds), *The Routledge Companion to Epistemology*. London: Routledge.
Hetherington, S. 1999. 'Knowing Failably.' *Journal of Philosophy*, 96: 565–87.
Kelp, C. 2009. 'Knowledge and Safety.' *Journal of Philosophical Research*, 34: 21–31.
Pritchard, D. 2008. 'Greco on Knowledge: Virtues, Contexts, Achievements.' *Philosophical Quarterly*, 58: 437–47.
——. 2012. 'Anti-luck Virtue Epistemology.' *Journal of Philosophy*, 109: 247–79.
Pritchard, D., Millar, A. and Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
Riggs, W. 2002. 'Reliability and the Value of Knowledge.' *Philosophy and Phenomenological Research*, 64: 79–96.
Sosa, E. 2007. *A Virtue Epistemology. Apt Belief and Reflective Knowledge*, Vol. 1. Oxford: Oxford University Press.
——. 2010. 'How Competence Matters in Epistemology.' *Philosophical Perspectives*, 24: 465–75.
——. 2011. *Knowing Full-Well*. Princeton, NJ: Princeton University Press.

CHRISTOPH KELP is currently working as BOF Research Professor at KU Leuven. He is principal investigator of a project on knowledge first virtue epistemology, which is funded by KU Leuven and the FWO (Research Foundation Flanders). His research focuses on a variety of areas in epistemology including epistemic goals and norms as well as the nature and value of knowledge, justified belief and understanding.

9 Thanks to Ernest Sosa and two anonymous referees for a number of helpful comments on this paper.

https://doi.org/10.1017/epi.2014.21 Published online by Cambridge University Press