

Hoops and Barns: a New Dilemma for Sosa

Christoph Kelp

Cameron Boulton

Fernando Broncano-Berrocal

Paul Dimmock

Harmen Ghijzen

Mona Simion

Abstract

This paper critically assesses Sosa's normative framework for performances as well as its application to epistemology. We first develop a problem for one of Sosa's central theses in the general theory of performance normativity according to which performances attain fully desirable status if and only if they are fully apt. More specifically, we argue that given Sosa's account of full aptness according to which a performance is fully apt only if safe from failure, this thesis can't be true. We then embark on a rescue mission on behalf of Sosa and work towards a weakened account of full aptness. The key idea is to countenance a distinction between negligible and non-negligible types of risk and to develop an account of full aptness according to which even performances that are endangered by risk can be fully apt, so long as the risk is of a negligible type. While this alternative account of full aptness solves the problem we developed for Sosa earlier on, there is also bad news for Sosa. When applied to epistemology, the envisaged treatment of barn façade cases as cases in which the agent falls short of fully apt belief will no longer work. We show that, as a result, Sosa faces a new version of a familiar dilemma for virtue epistemology. Either he construes full aptness as strong enough to get barn façade cases right in which case his view will run right into the problem we develop. Or else he construes full aptness as weak enough to avoid this problem but then he will not be able to deal with barn façade cases in the way envisaged.

Introduction

In his recent book *Judgement and Agency* Ernest Sosa develops a general account of the normativity of performances with an aim according to which the fully desirable status of any such performance is what he calls full aptness. He goes on to apply this account to the domain of epistemology and identifies human knowledge

with a particular normative standing of belief, to wit, fully apt belief. The resulting account of human knowledge is said to provide a particularly appealing solution to a number of venerable epistemological problems, including the problems posed by Gettier cases and clairvoyance cases.

This paper takes a closer look at Sosa's account of performance normativity and its application to epistemology. More specifically, we will first outline Sosa's general account of performance normativity (§1), its application to the case of belief (§2) and the proposed solution to one key epistemological problem, to wit, the Gettier problem (§3). §4 raises a problem for Sosa's account of full aptness. We then show that a minor and independently motivated modification to this account serves to avoid this problem (§5). At the same time, we argue that adopting it will come at a cost for Sosa as the proposed solution to the Gettier problem will no longer work (§6). Finally, §7 argues that, as a result, Sosa's view threatens to succumb to a familiar dilemma for virtue epistemological accounts of knowledge.

1 Performance Normativity

According to Sosa, any performance with an aim (henceforth just 'performance' for short) can be assessed along three dimensions: success, competence and aptness. A performance is successful, according to Sosa, if and only if it attains its aim. It is competent if and only if it is produced by the exercise of a competence to attain its aim and it is apt if and only if it is successful because competent. By way of example suppose you are an archer and you have just taken a shot which is aimed at hitting a certain target. According to Sosa's account, we can assess your shot along the three dimensions just mentioned. That is to say, we may ask (i) whether it is successful, i.e. whether it does hit the target, (ii) whether it was competent, i.e. whether it was produced by the exercise of a competence on your part to hit the target, and (iii) whether it was apt, i.e. whether it hit the target because of the exercise of your competence to hit it.

In addition, Sosa distinguishes among three types of aptness: animal aptness (a-aptness), reflective aptness (r-aptness) and full aptness (f-aptness).

A-aptness is simple first order aptness. That is to say, it is first order success because of first order competence. Returning by way of illustration to the archery case, your shot is a-apt if and only if it hits the target because of an exercise of your first order competence to hit the target.

Besides first order abilities normal human agents also have higher order abilities, including monitoring abilities that enable them to gauge the risk of inaptness of a given performance. These monitoring abilities are relevant to reflectively competent performance, which Sosa illustrates by means of the following example:

[An archer's shot] is reflectively competent if and only if it corresponds to a competent second order awareness¹ that the shot would

¹ Note that Sosa takes the relevant sense of 'awareness' to be *non-factive*. This is abundantly clear from the the passage right after the one quoted above: "Either of these—aptness and reflective

be apt.

[Sosa 2015: 68]

It is clear from Sosa's account that he takes this biconditional to generalise beyond the case of shots in archery, if not to performances in general so at least to an important subclass of performances. In what follows, we will assume that this generalisation does hold. (We will also take any further restriction to subclasses of performances that may be needed as read.)

Given that this is so, it is natural to assume that one's performance is reflectively *apt* if and only if it corresponds to an *apt* second order awareness that one's performance would be apt. Suppose you, the archer, take an apt shot at a certain target. Before taking the shot you assess the risk of inaptness: are lighting conditions sufficiently favourable, are winds sufficiently normal, are you sufficiently concentrated and sober and so on. Your assessing the risk of inaptness is itself a performance and may be assessed in terms of success, competence and aptness. The thought is that your shot will be r-apt if and only if it is not only a-apt but also accompanied by a risk assessment that outputs an apt awareness that your shot would be apt.

On this account of r-aptness, any r-apt performance will be *safe from inaptness* at the first order. Here is why. Since apt awareness is factive, one arrives at an apt awareness that one's performance would be apt only if one's performance would indeed be apt. But if one's performance would be apt it is in that way safe from inaptness. So, r-apt performance requires one's performance to be safe from inaptness on the first order.²

competence—can be present without the other.” [Sosa 2015: 68] It's easy to see that if reflective competence can be present without aptness, then the sense of 'awareness' at issue in the account of reflective competence cannot be factive. For the purposes of this paper, we will adopt Sosa's non-factive usage of 'awareness'.

² But doesn't Sosa explicitly deny that aptness requires safety? Even if at some point he did, there is reason to believe that he now does accept a safety condition on aptness. Consider the following passage, which comments on a case of a basketball player who takes a shot from barely within her range but isn't aware of this fact (this is effectively the case of player B from section 4 below):

Our basketball example suggests a distinction between first-order safety and second-order safety. The player's shot will be safe when she is (even barely) above her threshold of sufficient reliability. So situated, not too easily would she then fail in her attempt to make that goal. Unlike that first-order performance, however, her second-order performance may still be unsafe. Unaware of her threshold's location, she might too easily have shot inaptly, below that threshold. Properly situated as she is in fact (though barely), she is thereby disposed to shoot successfully and aptly. Because she is unaware of her threshold, however, she might too easily have shot inaptly. She might so easily have been improperly situated and still have shot just the same.

[Sosa 2015: 72]

What Sosa suggests here is that even a-aptness features a safety condition, although one that, in the basketball case, the player satisfies. Even if one may still want to dispute that Sosa accepts the safety condition on a-apt performance, there can be no question that he accepts a safety condition

Finally, in addition to a-aptness and r-aptness, full aptness requires a connection between the two. Here is Sosa once more:

A performance is fully apt if and only if it is guided to aptness through the agent's reflectively apt risk assessment.

[Sosa 2015: 69]

For instance, you may produce an apt shot that was aptly risk-assessed to be such that it would hit the target, without it being the case that your producing this shot has much to do with this assessment. Having arrived at an apt risk assessment that your shot would be successful, you may still may decide whether to take it by a coin toss. In that case, the aptness of your shot will not derive from the apt awareness that your shot would be apt. Even if a-apt and r-apt, your shot will not be f-apt.

It may also be worth noting that since f-aptness entails r-aptness, f-apt performance requires safety from inaptness in the same way as r-apt performance.

Now, here is one of Sosa's key theses:

The Normative Thesis (NT). "The fully desirable status for performances in general is full aptness: it is aptness on the first order guided by apt awareness on the second order that the first order performance would be apt (likely enough)." [Sosa 2015: 85]

Sosa also says that a performance "suffers" or "falls short" if it doesn't rise to the level of f-aptness and that f-aptness is "what really matters" [2015: 85]. The thought here is that f-aptness picks out a special standing of performances: f-apt performances are fully non-defectively and enjoy special value.

For the purposes of this paper, we will go along with Sosa and assume that NT is true.

2 Belief as an Epistemic Performance

Sosa also takes belief to be a type of performance, to wit, a type of epistemic performance that constitutively aims at truth. Given that this is so, the general normative framework can be applied to the particular case of belief.

On the first order, we can assess beliefs in accordance with (i) whether they attain their aim, i.e. whether they are true, (ii) whether they are competent, i.e. whether they are produced by an exercise of a competence to form true beliefs, and (iii) whether they are a-apt, i.e. whether they hit the mark of truth because of the exercise of a competence to form true beliefs.

However, we can also ask whether beliefs attain higher order kinds of aptness, besides a-aptness. In particular, we may ask whether a belief is r-apt, i.e. whether

on r- and f-apt performance. This is already abundantly clear in the above passage and further confirmed in his discussion of barn façade cases (see also below), where he explicitly ventures to explain the intuition of ignorance in terms of lack of safety [Sosa 2015: 79].

it is a-apt and an assessment of risk has led to an apt awareness that it would be apt. We may also ask whether it is f-apt, i.e. whether it is guided to aptness by r-aptness.

Crucially, Sosa proposes to identify human knowledge with a particular normative standing in his account of performance normativity as applied to belief, to wit, f-apt belief. In other words, he also endorses the following thesis:

The Epistemological Thesis (ET). Human knowledge is f-apt belief.³

For the purposes of this paper, we will also assume with Sosa that ET is true.

It may be worth noting that ET has a number of important benefits, especially in conjunction with NT. First, it promises to offer an attractive solution to the value problem, i.e. the problem of explaining why human knowledge is better than belief that falls short of human knowledge. By NT, only f-apt performances are fully non-defective. Since, in general, it's better to perform fully non-defectively than not to do so, by NT, we get the result that f-apt performance is better than performance that falls short of f-aptness. What goes for performances in general also goes for the special case of belief. As a result, f-apt belief is better than belief that falls short of f-aptness. Since, by ET, human knowledge is f-apt belief, human knowledge is better than belief that falls short of human knowledge. Second, as Sosa argues in some detail, ET offers an appealing solution to a number of key epistemological problems, including the Gettier problem and the problem posed by clairvoyance cases. Here we want to take a brief look at how Sosa ventures to deal with the Gettier problem, and, in particular, with the notorious barn façade case [Goldman 1976 who, in turn, attributes it to Carl Ginet].

³ This is not to say that Sosa doesn't also want to distinguish between animal and reflective knowledge, nor that these are kinds of knowledge that humans often attain. That said, some of Sosa's remarks are very much in line with ET, for instance when he describes his "account of human knowledge as *requiring knowledge full well*" [Sosa 2015: 85, *our emphasis*]. At the same time, not everything he says chimes equally well with ET. Consider, for instance, the following passage: "The account of a *desirable level* of human knowledge as knowing full well . . ." [Sosa 2015: 85, *our emphasis*]. There is at least an implicature here that there are other levels of human knowledge besides knowledge full well. If so, it is hard to see how human knowledge could still be identified with knowledge full well/f-apt belief. And perhaps this latter statement is the one that better captures what Sosa has in mind.

Fortunately, we don't have to settle this issue here. The reason for this is that the arguments below will go through no matter whether Sosa accepts ET or the alternative claim about what it takes to attain a desirable level of human knowledge. After all, it is clear that Sosa aims to explain the intuitive absence of knowledge in certain Gettier cases—notably, barn façade cases—in terms of the absence of f-apt belief. Whether he takes this to amount to an explanation in terms of the absence of human knowledge (as ET would have it) or to an explanation in terms of the absence of a desirable level of it (as the alternative claim would have it) is by-the-by, at least for present purposes. And while, in what follows, we will stick to ET, this is in the first instance to maximise readability.

3 Barn Façade Cases

Suppose you are driving down a road that is lined by fields. You see what looks like a barn to you and thereupon come to believe that the structure you are looking at is a barn. Since you are a reliable barn spotter, your belief is justified. Moreover, your belief is also true: the structure you are looking at is in fact a barn. Unbeknownst to you, however, it is the only real barn in an area that is otherwise peppered with mere barn façades that are indiscriminable from real barns from your position on the road.

Intuitively, your belief that the structure is a barn falls short of knowledge. One major stumbling stone for any virtue epistemology that ventures to identify knowledge with a-apt belief is that there is excellent reason to believe that your belief is a-apt in the sense that it is true because competent. After all, in parallel non-epistemic cases, the relevant performances are a-apt also. For instance, if you are at a shooting range and take a shot at the only target that hasn't been sabotaged to prevent any shot fired at it from hitting, your shot is still a-apt [Pritchard 2008].

Sosa's account of human knowledge promises to do better. The reason for this is, of course, that Sosa's account identifies human knowledge not with a-apt belief but with f-apt belief. Since, recall, f-apt performance in general requires safety from inaptness, so does f-apt belief. Crucially, this safety requirement on f-apt belief is not satisfied. After all, you are in conditions in which nearly nothing that looks like a barn is a barn. As a result, you might easily have ended up looking at a mere façade in which case you would have formed a false (and hence inapt) belief that the structure you are looking at is a barn. It is thus not the case that your belief that the structure you are looking at is a barn is safe from inaptness. Your belief that the structure you are looking at is a barn falls short of f-aptness. By ET, your belief falls short of human knowledge. The intuition of ignorance is duly accommodated and one of the major obstacles for virtue epistemology is overcome [Sosa 2015: 79].⁴

4 Sosa in Trouble

Let's take a closer look at Sosa's account of f-aptness. For starters, let's first look at a case that Sosa also uses in developing his account: two basketball players, A and B, take the same type of shot, one that is just within the range within which they have the competence to score reliably enough. A also has an apt awareness that his shot would be apt, and takes the shot guided by this apt awareness. In contrast, B does not have an apt awareness that his shot would be apt. The reason for this, let us suppose, is that B doesn't have a clear view on just how far his range

⁴ It may be worth noting that a considerable number of virtue epistemologists have taken on the problem of barn façade cases besides Sosa, including two of us (Broncano-Berrocal [Forthcominga, Forthcomingb], Kelp [2011, 2013, Forthcoming]), John Greco [e.g. 2010, 2012], Clayton Littlejohn [2014], Alan Millar [e.g. 2010], Duncan Pritchard [2010, 2012], and John Turri [e.g. Turri 2016]. However, since this issue is devoted specifically to the epistemology of Ernest Sosa, we will set alternative approaches to barn façade cases aside here.

extends. Let us also suppose that, as a result, B might just as easily have taken the shot from outside his range.

According to Sosa, between A and B, only A's shot is f-apt. After all, unlike A, B might very easily have taken a shot from outside his range. Since in that case his shot would not have been apt, B's shot does not satisfy the safety condition on f-apt performance. NT delivers the result that, between A and B, only A's shot attains fully desirable status.

This seems plausible enough. If this isn't immediately obvious, consider how a coach might react to the cases of A and B.⁵ In the case of B, he may fault the shot and criticise the player for taking a shot that, for all he knew, was out of his range. However, this would make little sense if B's shot had attained fully desirable status. So far, so good.

But now consider the following variation of the basketball case featuring A and C. C is just as competent as A (both at the first and at higher orders) and takes a shot from exactly the same position as A did. C also assessed the risk and arrived at an awareness that his shot would be apt and takes the shot guided by this awareness. In other words, the case of C is exactly like the case of A, with one crucial difference: there is an angry fan who threw a tomato at C which, let us assume, would have caused C to miss. However, his attempt was spoiled by a man in the row in front of him who happened to get up just when he threw the tomato and ended up taking the hit.

According to Sosa, between A and C, only A's shot is f-apt. After all, unlike A, C might very easily have been hit by a tomato. Since in that case his shot would not have been apt, B's shot does not satisfy the safety condition on f-apt performance. By NT, it follows that, between A and C, only A's shot attains fully desirable status.

Unlike in the case of A and B, here we would want to insist that this is the wrong result, for the following two reasons.

First, consider how a coach might react to the cases of A and C. In particular, let's ask whether the coach may fault C's shot just as in the case of B? This doesn't seem plausible. True, the coach may plausibly enough expect his players to have some sense of how far their ranges extend. He may plausibly enough criticise players that don't and try to get them to acquire such a sense. That's why it is plausible that B's shot falls short of fully satisfactory status. Crucially, however, the same is not true in the case of C. It is not as if the coach may plausibly enough expect his players to have some sense of whether they are threatened by angry tomato-throwing fans. Nor may he plausibly enough criticise players that are under such attack or get them to acquire a sense for when they are. There is no parallel reason to think that C's shot falls short of fully satisfactory status. In fact, if a player turned around before each shot to check whether there are tomato throwers among the spectators, the coach could appropriately criticise him for *that*. C did well, then, to go ahead and take the shot rather than first check whether he is under

⁵ Note that the idea of thinking about how a coach might react is due to Sosa himself [Sosa 2015: 71], albeit in the context of a slightly different kind of case.

attack by angry fans. In this way, there is reason to think that not monitoring for tomato throwers did not in any way tarnish C's shot. Since C's shot is just like A's in all other respects, there is reason to think that it attains fully satisfactory status if A's does.

Second, suppose that the coach is charged with giving an award for the best shot in the game. Suppose it's down, respectively, to A's shot vs. C's. Even if, in the earlier case of A and B, it makes sense for the coach to award A's shot over B's, the same does not hold when A's shot is compared to C's. Here it would be rather unreasonable for the coach to award the prize to A over C. If this isn't immediately obvious, suppose that after the coach declared the prize shared between A and C, A complains that C's shot was importantly inferior because there was an angry fan who might very easily have prevented C's shot from scoring by hitting him with a tomato. In this case the coach may legitimately ask A whether he is out of his mind and stick to his verdict that the prize be shared between A and C. However, this is very hard to reconcile with the idea that, between A and C, only A's shot attains fully desirable status.

What comes to light, then, is that performances can attain fully desirable status even though they are not safe from inaptness. Since we are granting Sosa NT, this means that he must be wrong about the safety condition on f-apt performances. Fortunately, as we will argue momentarily, the problem Sosa faces here is fixable.

5 Rescue Mission

Let's ask what's going on in the above cases. In particular, between B and C, why would only C's performance attain fully satisfactory status, especially since, in both cases, the performances are afflicted by a very high risk of failure? Here is one attractive answer. The risk C is running—being targeted by a tomato-throwing fan—is in a relevant sense *negligible*, whereas the risk B is running—taking a shot from outside his range—isn't. And while non-negligible risks tarnish performances, negligible ones don't.⁶ In particular, being afflicted by a negligible risk does not prevent a performance from rising to fully satisfactory status, whereas non-negligible ones do. That's why B's performance, which is afflicted by a non-negligible risk, falls short of fully satisfactory status, whereas C's performance, which is afflicted only by a negligible risk, doesn't.⁷

Since we are still granting Sosa NT—that performances attain fully desirable

⁶ Unless, of course, they materialise and actually spoil the performance. However, in what follows, we will simply take it as read that this is not what's happening.

⁷ Note that this may not be the only way of fixing the problem for Sosa. Fortunately, this is fairly unproblematic, at least for the purposes of this paper. The reason for this is our main critical point against Sosa—the dilemma we will outline in section 7—will stand no matter how the problem is fixed. All we need is that the problem is indeed fixed in the sense that, between B and C, C's performance comes out as rising to fully satisfactory status whereas B's doesn't. (We'll explain why in fn.11 below.) What we are interested in achieving here is to show that there is at least one way of fixing this problem that is compatible with what Sosa himself says about risks and to spell out in some detail how this can be done.

status if and only if they are f-apt—this means that we will need to modify the account of f-aptness. In particular, we will have to exempt performances afflicted by negligible risks from falling short of f-aptness. Here is our proposal: one’s performance is f-apt if and only if it is guided to aptness by an apt awareness that the aptness of one’s performance is not endangered by *non-negligible* risk. Since the risk B is running is non-negligible, his performance ends up falling short of f-aptness. Since the only risk afflicting C’s performance is negligible, his performance does rise to the level of f-aptness.

It might be thought that the difference between non-negligible and negligible risks is one that Sosa himself countenances. To see this note that Sosa does not hold that in order to attain fully satisfactory status, a performance cannot be afflicted by any degree of risk whatsoever. Rather, he only claims that the degree of risk cannot be too high. More specifically, Sosa holds that there is a “threshold below which the agent’s attempt, given their SSS situation, would be too risky” [Sosa 2015: 73]. The thought here is that so long as the degree of risk afflicting one’s performance is low enough to remain above the threshold, it is negligible, whereas once it increases to the point where it drops below the threshold, one is running a non-negligible risk.

We think that Sosa is right in distinguishing between negligible and non-negligible risks in this way. The trouble is that this just won’t serve to explain what we want to explain, i.e. the difference between B and C. To see this, note that the distinction is cashed out as a *quantitative distinction*: whether or not one is running a non-negligible risk will depend on whether the quantity of risk afflicting one’s performance is below a certain threshold. But of course, in the cases of B and C, the degree of risk they are running is the same (or so we may assume). As a result, if one drops below the threshold and so is running a non-negligible risk, so does the other. (Note also that Sosa will do well to allow that both B and C remain above the threshold in this particular case. Otherwise he simply won’t be able to secure the result that C’s performance attains fully satisfactory status.)

Our proposal is to countenance a distinction between negligible and non-negligible *types of risk* in addition to the distinction between negligible and non-negligible *amounts of risk*. In other words, we propose to countenance a *qualitative distinction* in addition to the quantitative distinction we can effectively find already in Sosa. The thought would then be that being targeted by an angry fan is a negligible type of risk. In contrast, taking a shot from outside one’s range is a non-negligible type of risk. This would give us just what we need. C is in all respects like A with the exception that C is at risk of failure. However, the risk is of a negligible type as a result of which it does not prevent C’s performance from rising to fully satisfactory status. B, on the other hand, is running a non-negligible risk, which means that his performance is bound to fall short.

While it may be plausible enough that the risk type of being targeted by an angry fan is negligible, whereas the risk type of taking a shot from outside one’s range isn’t, we may wonder whether more can be said about what exactly sets negligible types of risk apart from non-negligible ones. Sosa takes up the parallel

question for the quantitative distinction in the following passage:

What sets such a threshold? This will vary from domain to domain. It may be conventional and formalized, as in some professional contexts, or it may be less formal, more intuitive, as in the domain of a hunt. In each case, the threshold will be set by considerations distinctive of the domain and the proper basic aims of performances in it ...

[Sosa 2015: 73]

We can of course offer a parallel answer to the question at hand. Whether or not a certain type of risk is negligible is determined by domain-internal considerations such as the proper basic aims of performance types in it. In the domain of basketball the risk type of being targeted by an angry fan is negligible for the performance type shot, whereas the risk type of taking a shot from outside one's range isn't.

At the same time, some might be left with at least a slight feeling that this response isn't fully satisfactory. After all, it doesn't give us any indication about precisely how considerations about domains and basic aims of performances fix the threshold for reliability or which types of risks are negligible for that matter. In this way, one couldn't exactly say that it was maximally informative. But perhaps it's the best we can do. On the other hand, if we can do better than this, those who feel unsatisfied have a point.

Fortunately, a more informative answer can be given, at least in the case of negligible vs. non-negligible types of risk. To see how, recall that, according to the proposal we are considering, performances attain fully desirable status if and only if they are *f*-apt, where *f*-aptness requires an apt awareness that one's performance is free of certain types of risk. In other words, in order to rise to the level of *f*-aptness, one must have aptly monitored for certain types of risk. The question we'd like to ask now is: given all this, what types of risk does it make sense to monitor for in the first place?

To answer this question, we'd like to first sketch an (admittedly somewhat idealised) informal model of performing agents. Agents produce performances. They can be more or less reliable at producing successful performances of a certain type. All else equal (including degree of productivity (see below)), agents are better at producing performances of a certain type if they are more reliable. Agents can also be more or less productive at producing performances. All else equal (including the degree of reliability), agents are better at producing performances of a certain type if they are more productive.

While reliability and productivity are both good-making features of performing agents, there isn't a clear hierarchy between the two. To see this, consider two agents, D and E. Suppose that while D is more reliable than E at producing performances of a certain type, P, the difference is ever so slight. At the same time, E is significantly more productive than D. In this case, E may well be better at P than D. In that case, the benefits of productivity outweigh those of reliability. On the other hand, it might also be that E is ever so slightly more productive than D at P. At the same time, D is considerably more reliable than E. Here the benefits

of reliability may well outweigh those of productivity. Without trying to settle the difficult question of how to weigh reliability and productivity, let's agree that agents are better or worse at producing performances of a certain type, depending on whether they achieve a better or worse *balance between productivity and reliability* (henceforth BALANCE for short).

With this picture in play, let's return to the question of what types of risk it makes sense to monitor for, given that f-aptness requires apt monitoring for certain types of risk and that all and only f-apt performances attain fully desirable status. First, we want to say that it makes sense to monitor for a certain type of risk only if so doing improves agents' BALANCE. The reason for this not hard to fathom. If monitoring for a certain type of risk doesn't make one better at producing performances of the type in question, it makes little sense to monitor in the first place.

While this is an important part of our answer, it is not the full story. To see this, note that monitoring comes at a cost. At the very least, agents will incur the cost of acquiring the relevant monitoring equipment and the cost of running it. What we want to suggest is that, in addition, it makes sense to monitor for a certain type of risk only if the benefits of improved BALANCE isn't outweighed by the costs associated with monitoring. Again, the reason for this is easy enough to see. If the costs of monitoring outweigh the benefits, it doesn't make sense to go to the trouble of monitoring.

Finally, we want to suggest that these two necessary conditions are also jointly sufficient (at least in the idealised model we are considering). The thought then is that it makes sense to monitor for a certain type of risk if and only if (i) so doing improves agents' BALANCE and (ii) the benefits of improved BALANCE isn't outweighed by the costs incurred by going to the trouble of monitoring. We can then say that non-negligible types of risk are those types of risk that it makes sense to monitor for in this sense, and negligible ones are those that it does not make sense to monitor for.⁸

There is reason to think that this way of cashing out the distinction between negligible and non-negligible types of risks will deliver the desired results in the cases of B and C. Let's start with the case of B. First, having a sense of how far one's range extends and exercising it promises an increase in reliability as the risk of taking a shot from outside one's range is omnipresent in basketball. What's more, the increase in reliability very plausibly outweighs the costs in productivity.

⁸ We'd like to emphasise that we do not mean to say that the value of monitoring is exhausted by its BALANCE-enhancing potential. All we really need is that this is one good thing about monitoring. (Strictly speaking, then, when we say things like 'it makes sense to monitor', these claims should really be read as relative to the specific value associated with improvements in BALANCE.) After all, what we are trying to do is to offer a more detailed account of the distinction between negligible and non-negligible risks. The above account is one way to achieve this. It's compatible with all this (i) that there may be other values attached to monitoring, (ii) that these other values also serve to flesh out the difference between these two types of risk and (iii) that the alternative accounts are better than the one we developed above. We do not mean to take a stance on this here. If there is a better story than ours, we'd be the first ones to be interested in hearing it.

After all, alternative courses of action that might also lead to success (passing, continuing to dribble) tend to remain available in basketball. In this way, it is plausible that monitoring does improve agents' BALANCE. In addition, it seems also plausible enough that the benefits of this increase in BALANCE outweigh the costs of acquiring the sense of how far one's range extends and of exercising it. It's not particularly difficult to acquire at least some such sense simply by keeping a tab on past successes and failures. As a result, the risk of taking a shot from outside one's range is a risk it makes sense to monitor for and so comes out as non-negligible for the performance type basketball shots.

Let's move on to the case of C then. Does monitoring for tomato throwers make one a more reliable shooter? Since the threat of tomato throwers materialises ever so rarely (at both the actual world and at nearby worlds), if there is any gain worth mentioning at all here, it's going to be pretty minimal. What about productivity? Well, it would seem that, in the case of basketball, productivity threatens to be compromised, and arguably considerably so in comparison. After all, the time that one takes to check for tomato throwers is time that one loses for taking shots. If so, we may expect that monitoring for tomato throwers will mean a decrease in BALANCE with the result that this type of risk is negligible in basketball.

But suppose we grant that (i) checking for tomato throwers is neutral with respect to productivity and (ii) as a result, there is an overall improvement to the BALANCE. Suppose we also grant that, throughout your long and successful career in basketball, you are targeted once by a tomato thrower and so the type of risk we are considering does really materialise for you. (We assume that this would still place you well above the average for basketball players.) Even so, if you turned around before each shot to check for tomato throwers, your coach would be well advised to try and get you to refrain from doing this. In fact, it would be hard to categorise you in any other way than as paranoid. We take this to be pretty solid evidence that the costs of monitoring for tomato throwers outweigh the benefits of the improvement in BALANCE. And of course, this means that the second condition on non-negligible types of risk is not satisfied anyway. As a result, the risk of being targeted by a tomato thrower is a risk it does not make sense to monitor for and so comes out as negligible for basketball shots.

To sum up, Sosa's account of f-aptness came under pressure from the cases B and C. However, what has come to light now is that a fairly minor and independently plausible modification of the account, which restricts the condition of apt risk assessment to non-negligible types of risk, serves to avoid the problem. What's more, we have even been able to give an attractive account of the distinction between negligible and non-negligible risks that allows us to secure the desired results about the cases. We take all of this to be good news for Sosa.

6 Barn Façades Again

Unfortunately, not all is well for Sosa in the end. To see why not, we'll have to return once more to barn façade cases. Recall that agents in barn façade cases

intuitively do not know. While virtue epistemologists have struggled to accommodate this intuition, Sosa appeared to be able to do so in a well-motivated and elegant manner. Here it is again: By ET, human knowledge is f-apt belief. F-apt performances in general require safety from inaptness. So, f-apt belief requires safety from inaptness. The beliefs of agents in barn façade cases are not safe from inaptness. So, their beliefs fall short of f-aptness and do not qualify as human knowledge.

The problem is, of course, that §4 provided reason to think that f-apt performance does not require safety from inaptness, at least given that we grant NT. By the same token, there is reason to believe that Sosa's account of the intuition of ignorance in barn façade cases remains unsuccessful, at least in its present form.

At the same time, we also provided an alternative account of f-aptness that can steer clear of the problem discussed in §4. Recall that, according to this account, one's performance is f-apt if and only if it is guided to aptness by an apt awareness that the aptness of one's performance is not endangered by non-negligible risk. On this account, a performance can be f-apt even when its aptness is endangered by risk. As a result, it does not give us the kind of safety condition on f-apt performance we get from Sosa's account of f-aptness. That said, it remains the case that performances the aptness of which is endangered by non-negligible risk are bound to fall short of f-aptness. This means that we still get some safety condition on f-apt performance, i.e. a safety condition that is restricted to non-negligible risk.

The question that naturally arises at this stage is whether Sosa's account of barn façade cases will work even once we have replaced his account of f-aptness by our weaker alternative. It is easy to see that this question will boil down to the question of whether the risk of looking at a mere barn façade is negligible for the performance type of forming perceptual beliefs about the presence of barns.

Unfortunately for Sosa, there is excellent reason to think that the answer here is that the risk is negligible. Recall that, a type of risk is non-negligible in the relevant sense if and only if (i) monitoring improves agents' BALANCE and (ii) the benefits of improved BALANCE isn't outweighed by the costs incurred by going to the trouble of monitoring.

Now, just as in the case of the tomato thrower, the risk of looking at a mere barn façade materialises ever so rarely (at both the actual and nearby possible worlds). As a result, the increase in reliability that comes with monitoring for this type of risk is marginal. At the same time, going to the trouble of walking all the way around each barn to check that it isn't a mere façade is a very time-consuming affair, which will likely affect one's productivity negatively and considerably so, at least in comparison with the gains in reliability. In that case, monitoring for mere façades will mean a decrease in BALANCE with the result that this type of risk is negligible for the performance type of forming perceptual beliefs about the presence of barns.

But now suppose we were to grant that monitoring for mere façades is neutral with regard to productivity and, by the same token, that there is an overall improvement to the BALANCE. Suppose we also grant that, throughout your long

and successful career as a barn spotter, you once spot a mere barn façade. (We assume that this would still place you well above the average for a barn spotter.) Even so, if you walked around each barn to check whether it isn't a façade, it would be hard to categorise you in any other way than as paranoid. Just as in the case of the tomato thrower earlier, we take this to be pretty solid evidence that the costs of monitoring for mere façades outweigh the benefits of the improvement in BALANCE. As a result, the risk of looking at a mere barn façade comes out as negligible for the performance type of forming perceptual beliefs about the presence of barns anyway.

Of course, if the risk is negligible, then, on the alternative weaker account of f-aptness we proposed above, the beliefs of agents in barn façade cases will come out f-apt.⁹ Given that ET holds, it follows that the weaker account predicts that agents in barn façade cases do attain human knowledge. As a result, the intuition of ignorance cannot be explained in the way envisaged.¹⁰

7 A Familiar Dilemma

In fact, the above considerations provide reason to think that Sosa is threatened by a version of a familiar type of dilemma that virtue epistemological accounts of knowledge face. In its standard form, the dilemma looks like this: if a candidate virtue epistemological account of knowledge is strong enough to successfully explain the absence of knowledge in certain Gettier cases (e.g. barn façade cases), then it is too strong to accommodate the presence of knowledge in cases in which knowledge is clearly present (e.g. cases of testimonial knowledge). And if the account is weak enough to accommodate the presence of knowledge in the latter cases, then it is too weak to explain the absence of knowledge in the former [e.g. Lackey 2007, 2009, Pritchard 2008, Pritchard *et al.* 2010, Broncano-Berrocal Forthcomingb, Kelp 2016a, Forthcoming].

The present version of the dilemma is slightly more complex. Here goes: if Sosa's account is strong enough to be able to successfully explain the absence of knowledge in certain Gettier cases (notably barn façade cases), he is bound to run into trouble with some other (non-epistemic) cases, to wit, the case of basketball player C. (This is what happens with Sosa's original account of f-aptness.) If the

⁹ If this isn't immediately obvious, note that just as, in the earlier basketball cases, the only difference between A and C is that C is running the risk of being targeted by a tomato thrower, the only difference between an agent in normal barn county and in barn façade county is that the latter is running the risk of looking at a mere façade. As a result, just as, in the basketball cases, if A's performance is f-apt, then so is C's, so, in the barns cases, if the agent in normal barn county arrives at an f-apt belief, then so does the agent in barn façade county.

¹⁰ Couldn't Sosa venture to avoid this problem by opting for a different account of negligible/non-negligible types of risk according to which the risk of looking at a mere barn façade comes out as non-negligible, at least when you are in barn façade county? Perhaps. Even so, the problem is that this will just get him right back into trouble with the case of C as we will argue in more detail momentarily (see especially fn.12).

account is weakened to avoid the difficulty posed by the relevant non-epistemic cases, it will be too weak to explain the absence of knowledge in the relevant Gettier cases. (This is what happens when the account of f-aptness is weakened along the lines suggested above.) Either way, the account is ultimately less than fully satisfactory.^{11,12}

Conclusion

This paper has put Sosa's account of the normativity of performances and its application to epistemology under scrutiny. In particular, we have focused on Sosa's account of f-aptness. We have argued that the account of f-aptness is too strong, at least given that Sosa is right about NT, i.e. the claim that f-aptness is the fully desirable status of performances. More specifically, we have shown that the safety condition that, on Sosa's account, f-apt performances must satisfy is too strong.

On a more constructive note, we have tried to fix this problem by offering a weaker account of f-aptness. The key move was to distinguish between negligible and non-negligible types of risk and to modify the account of f-aptness such that f-apt performances must be safe only from non-negligible risks. We have also worked toward a more detailed account of how to draw the distinction between negligible and non-negligible types of risk.

While our alternative does indeed fix the problem Sosa encounters, we have also argued that adopting it comes at a cost for Sosa. Sosa's takes his account of the normativity of performances to have important epistemological applications. In

¹¹ Recall that in fn.7 we said that, for the purposes of this paper, it doesn't matter how the problem posed by the cases of B and C is fixed. There is reason to think that main critical point against Sosa will stand so long as the problem is indeed fixed in the sense that, between B and C, C's performance rises to fully satisfactory status whereas B's doesn't. We can now see why. If C's performance is fully satisfactory, by NT, it is f-apt. But since C's performance is unsafe, it follows that f-aptness does not require safety. In this way, fixing the problem posed by the basketball case will commit Sosa to giving up the very condition on f-aptness that was meant to explain the absence of knowledge in barn façade cases. While this doesn't entail that Sosa will be unable to accommodate the absence of knowledge here, it is independently plausible that the case of C is a non-epistemic analogue of the barn façade case. As a result it is also plausible that C's performance rises to fully satisfactory status if and only if the beliefs of agents in barn façade cases do. Since this is all that we need to force the dilemma on Sosa, there is reason to think that Sosa will face the dilemma no matter whether the specific solution to the problem posed by B and C is the right one.

¹² Now we can also see why Sosa will not be able to avoid the problem from section 6 along the lines suggested in fn.10, i.e. by opting for an account of negligible/non-negligible types of risk according to which the risk of looking at a mere barn façade case is non-negligible, at least when you are in barn façade county. Given that the case of C is a non-epistemic analogue of the barn façade case, any account of negligible/non-negligible risk that will deliver the result that risk of looking at a mere barn façade is non-negligible at least when you are in barn façade county, will also deliver the result that the risk of missing a shot due to a tomato attack C is running is non-negligible also at least when C is under attack by a tomato throwing fan. Even if this move will enable Sosa to secure the desirable result that, in the barn façade case, your belief falls short of f-aptness (and thereby of human knowledge), it will firmly commit him to the undesirable result that C's shot falls short of f-aptness also (and thereby falls short of fully satisfactory status).

particular, he accepts ET, i.e. the claim that human knowledge is f-apt belief, and uses it to address a number of long-standing epistemological problems, including the Gettier problem. What has come to light is that, once the account of f-aptness is weakened in the way we proposed, the resulting version of ET is too weak to offer a satisfactory solution to the Gettier problem. In particular, it will no longer be able to accommodate the intuition of ignorance in barn façade cases.

In this way Sosa thus faces a familiar dilemma. Either he construes f-aptness in a strong enough manner to successfully solve the Gettier problem. In that case, however, he will run right into the problem posed by the basketball case. Or else he opts for an account of f-aptness that is weak enough to avoid this problem. But then it is too weak to solve the Gettier problem.

Of course, we have throughout assumed with Sosa that both NT and ET are true. So, for all we have argued, Sosa might still respond to the argument by recanting either (or both) of NT and ET. However, there are costs associated with this strategy also. If he gives up his account of knowledge as f-apt belief, he will have to start over on epistemological problems this account was meant to solve. Moreover, if he wants to hold on to the idea that fully desirable performances are f-apt performances, he will have to explain to us why we epistemologists should continue to be so concerned with knowledge rather than shifting our attention to f-apt performances. What if Sosa opts for holding on to the account of knowledge as f-apt belief and gives up the claim that performances attain fully satisfactory status if and only if they are f-apt instead? In that case, he faces a similar problem. It's just that now the question is why we epistemologists should continue to be concerned with knowledge/f-apt belief rather than with whatever it is that makes for the fully desirable status of belief. And if Sosa drops both claims, then he is back to the drawing board not only on the epistemological front but also on the normative one.

For what it's worth, we think that the way forward is to embrace the second horn of the dilemma, i.e. to opt for a weaker account of f-aptness that avoids the problems in cases of non-epistemic performances. We also think that a promising way of addressing the remaining epistemological problem is by requiring that the epistemic success at issue in apt belief (or perhaps in apt perceptual belief) is something stronger than mere truth. While one of us has explored the prospects of a number of proposals in this direction elsewhere¹³, we will not take up the issue in any more detail here.

References

BRONCANO-BERROCAL, F. Forthcominga. Purifying impure virtue epistemology. *Philosophical Studies*.

¹³ [Kelp 2011], which takes the relevant epistemic success to be discrimination, and [Kelp Forthcoming, 2016b], takes it to be knowledge.

- BRONCANO-BERROCAL, F. Forthcomingb. A robust enough virtue epistemology. *Synthese*.
- GOLDMAN, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy*, **73**, 771–91.
- GRECO, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- GRECO, J. 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research*, **85**, 1–26.
- KELP, C. 2011. In defence of virtue epistemology. *Synthese*, **179**, 409–33.
- KELP, C. 2013. Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, **91**, 265–78.
- KELP, C. 2016a. Epistemic Frankfurt cases revisited. *American Philosophical Quarterly*, **53**, 27–37.
- KELP, C. 2016b. Justified belief: knowledge first-style. *Philosophy and Phenomenological Research*, **93**, 79–100.
- KELP, C. Forthcoming. Knowledge first virtue epistemology. In: CARTER, A., GORDON, E., & JARVIS, B. (eds), *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- LACKEY, J. 2007. Why we don't deserve credit for everything we know. *Synthese*, **158**, 345–61.
- LACKEY, J. 2009. Knowledge and credit. *Philosophical Studies*, **142**, 27–42.
- LITTLEJOHN, C. 2014. Fake barns and false dilemmas. *Episteme*, **11**, 369–389.
- MILLAR, A. 2010. Perceptual knowledge and recognitional abilities. In: PRITCHARD, D., MILLAR, A., & HADDOCK, A. (eds), *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
- PRITCHARD, D. 2008. Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, **58**, 437–47.
- PRITCHARD, D. 2010. Relevant alternatives, perceptual knowledge and discrimination. *Noûs*, **44**, 245–68.
- PRITCHARD, D. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy*, **109**, 247–79.
- PRITCHARD, D., MILLAR, A., & HADDOCK, A. 2010. *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
- SOSA, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.

TURRI, J. 2016. Knowledge as achievement, more or less. *In*: FERNANDEZ, M. (ed), *Performance Epistemology*. Oxford: Oxford University Press.