

Sosa on knowledge, judgment and guessing

J. Adam Carter¹

Received: 22 April 2016 / Accepted: 23 July 2016 / Published online: 5 August 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract In Chapter 3 of *Judgment and Agency*, Sosa (Judgment and Agency, 2015) explicates the concept of a *fully apt* performance. In the course of doing so, he draws from illustrative examples of practical performances and applies lessons drawn to the case of cognitive performances, and in particular, to the cognitive performance of judging. Sosa's examples in the practical sphere are rich and instructive. But there is, I will argue, an interesting disanalogy between the practical and cognitive examples he relies on. Ultimately, I think the source of the disanalogy is a problematic picture of the cognitive performance of *guessing* and its connection to knowledge and defeat. Once this critical line of argument is advanced, an alternative picture of guessing, *qua* cognitive performance, is articulated, one which avoids the problems discussed, and yet remains compatible with Sosa's broader framework.

Keywords Ernest Sosa · Virtue epistemology · Knowledge · Performance epistemology

1 Introduction

The topic of this paper is a particular kind of cognitive performance—*guessing*. As Sosa (2015) sees it, guessing is like judging, but not quite. Both guessing and judging have an alethic aim. When we guess, as well as when we judge, we aim to get it right. But judgment, on Sosa's view, does not aim *merely* to get it right. Judgment has a further epistemic aim, which is to get it right aptly. And getting it right aptly involves competent risk assessment.

✉ J. Adam Carter
j.adam.carter@ed.ac.uk

¹ University of Edinburgh, Edinburgh, UK

Competent risk assessment—and not just hitting the relevant target any old way—is obviously important when engaged in sports, such as basketball. Suppose you shoot from *far* beyond your threshold of reliable competence—by chucking the ball toward the basket from the half court line. This (assuming the shot clock is not expiring) is very poor shot selection in the context of a basketball game. Your coach and teammates will rightly berate you for taking such a foolish shot, even if it happens to go in. Things are very similar in the cognitive arena: even if a wild guess turns out to be right by happenstance—and even though (like the shot taken from half court) you *aimed* by guessing to get it right, something is lacking. A correct but wild guess (like the shot from half court that lands in the basket) is not *apt*, which is to say, its accuracy does not manifest the shooter’s competence to make the shot reliably enough.¹

Of course, half-court shots in basketball games are *very* rarely ever taken. Likewise, in inquiring, we very rarely—unless practically forced²—aim to get it right by means of a wild guess.³ However, compare shooting from half-court (beyond one’s threshold for reliability) with shooting *just barely within* one’s threshold for reliability, though in a circumstance where you can’t tell that you are in fact within this threshold. This case is clearly *different* from the case where you shot from half-court. Your shot in this case is *apt*, accurate because adroit. Yet, there is still a kind of credit-reducing luck which taints the assessment of the performance. Put simply, for all you knew, when you shot, you were outside your threshold of reliability. You just happened to be within this threshold. More carefully: your performance was *first-order* safe, but not *second-order* safe.⁴ You are just reliable enough from *that* distance that not too easily would you fail to make your shot from there. But because you are unaware of where this threshold lies, you too easily could have shot *inaptly*, from just beyond your threshold of sufficient reliability.

Just as a basketball shot, under these circumstances, is not fully apt, neither is what Sosa takes to be the *cognitive analogue* to such a basketball shot: a guess under conditions under which you *are* just above the threshold for sufficient reliability, *but aren’t aware of this*. Such a guess (unlike a wild guess) would be first-order safe. But it would not be second-order safe because, unaware of where the threshold lies, you

¹ The matter of how reliable is ‘reliable enough’ is not invariant across domains of endeavour but is sensitive to a range of factors specific to particular domains. As Sosa (2015) puts it, ‘In basketball we know at least roughly where it lies, with due allowance for the position of the player and his teammates, the time remaining on the clock, whether the shot to be taken is a three-point shot, etc. Many factors thus bear in diverse ways, and good players will take them into account, aiming not just to sink a basket, but to manifest in so doing the full competence required’ (2015, p. 71). See also Sosa (2015, Ch. 8), especially pp. 168–180 for a more detailed discussion of what it takes to be reliable enough, situated within the context of pragmatic encroachment.

² This would be the cognitive analogue to being practically forced to shoot, as when the shot clock is expiring. An example of such a cognitive situation might face a traveler, starving and thirsty, who must choose which way to go at a fork in the road, only one road she is told, leads to food and water. Here, practical circumstances force the traveler to aim to get it right by taking a wild guess. But, were the traveler not under such duress, the traveler would aim to get it right, but not by taking a wild guess—rather, by a more reliable means.

³ Wild guessing, under no kind of duress, is poor epistemic shot selection.

⁴ See Sosa (2015, p. 72).

too easily could have guessed inaptly. Such a guess would be apt, but not *fully* apt; it would aspire to animal, but not reflective, knowledge.⁵

The foregoing captures only in very broad brush strokes Sosa's thinking about guessing, and how guessing relates to other notions within his virtue epistemology. He has many more illuminating things to say. In fact, his remarks on the topic of guessing (and its analogues, in the case of practical performances) crop up at various places throughout *Judgment and Agency*. In particular, what Sosa says about guessing plays an important and focused role in his explication in Chapter 3 of the concept—central to *Judgment and Agency*—of a *fully apt performance*. This distinguishes guessing as important to Sosa's wider project in the book, given that he applies the concept of a fully apt performance in his account of the nature and normativity of human judgment—the latter being his primary aim.

In this paper, I offer a focused critique of Sosa's discussion of guessing, with particular attention paid to the role which guessing plays in Sosa's dialectical project in Chapter 3, and to a lesser extent Chapters 6 and 8 of *Judgment and Agency*. In response to the objections raised, I offer an alternative way of thinking about guessing which avoids these objections and yet which remains compatible with Sosa's wider virtue epistemological framework.

Here is the plan. Section 2 details Sosa's account of guessing, taking as a starting point his EYE EXAM case, which is an important reference point in Chapter 3. Section 3 introduces the notion of *performative defeat* and highlights, with reference to performative defeaters, some structural differences between EYE EXAM and other cases Sosa presents as on a performative par with EYE EXAM. Section 4 draws from the Jamesian insight that is that the aim of possessing the truth is best understood as a mixture of *two competing aims*—truly representing and not misrepresenting—to suggest that guesses, including reliable ones, generate defeaters in a way that (contra Sosa) disqualifies them from *knowledge*, animal or otherwise. Section 5 discusses the approach to guessing sketched thus far in the wider context of unreliable achievements, and distinguishes various species of guesses in connection with what are presented as practical performative analogues.

2 Sosa on guessing

Sosa (2015) gives the following example:

EYE EXAM: When I go for my exam, I am asked to read the lines of a chart with letters that shrink line by line from a huge single letter at the top, to those barely visible at the bottom. At some point I start to lose confidence that I am getting the letters right, but I keep going until the technician tells me to stop and then records some result. At that point there are many cases where I am quite unsure as to whether it is an "E" or an "F", say, or a "P" rather than an "F", etc. Suppose, however, it turns out that (unbeknownst to me) I am in fact unfailingly right year after year at a line where I am thus unsure. At that point I am in effect

⁵ See also Sosa (2010) for discussion on the difference between first-order and reflective competence.

“guessing.” I do affirm, to myself in private and to the technician in public, and I do so in the endeavor to get it right. That is after all what the test requires: that I try to answer correctly. And we can surely stipulate that I thereby manifest a competence, one I do not recognize as reliable enough. *This latter* is why I resort to guessing, when I continue to affirm as I undergo the test. Unbeknownst to me, however, my affirmations turn out to be surprisingly reliable (2015, pp. 74–5).

EYE EXAM—specifically, the circumstance where one reliably reads the last line, while unaware that one is reliable in doing so—features what we can call *reliable guessing*. It’s reliable because the examinee is unfailingly right in affirming the final line in the eye test. But, as Sosa tells us, it’s nonetheless guessing; you don’t regard your competence as reliable enough in the case of reading the last line, even though you are.

Here are three central claims Sosa makes about reliable guessing. Firstly, reliable guessing—like guessing more generally—has a (merely) *alethic aim*; in EYE EXAM, the agent aims by guessing to get it right, but not to get it right *aptly*.⁶

Second, the reliable guesser’s belief is a kind of *subcredal animal knowledge* (2005, p. 76): his guesses are not right by luck (Ibid., p. 76); rather, the guesser affirms in a way that manifests a reliable first-order visual-perceptual competence which is however not recognised by the guesser as reliable enough.⁷

Thirdly, what distinguishes the knowledge the examinee enjoys of the large *top-row* letters from the kind of (subcredal) animal knowledge she enjoys of the small, bottom row in virtue of, in the latter case, reliably guessing? Sosa’s answer here is nuanced. He first considers, and then, rejects, that the epistemically relevant difference between these two cases of affirming with the aim of getting it right is just that the examinee affirms with comparatively more *confidence* the higher-row letters than the bottom-row letters (even though that is of course the case).⁸

Rather, Sosa thinks, the difference in epistemic status is to be explained specifically in terms of a difference in reflective, or second-order, competence in the two cases. His remarks on this point are illuminating:

⁶ As Sosa (2015, p. 75) puts it, ‘We still affirm at the lower rows, while aiming *to get it right*. We give it our best shot, since only thus will we undergo the vision test properly. But do we aim to get it right, *aptly, reliably enough*? No, by that point it matters little whether we do get it right *reliably at all* [...] we make our guess without *endeavouring* to get it right aptly’. See also Sosa’s fn. 14.

⁷ See also Sosa’s remarks in Chapter 6, following his discussion of the case of Simone. Here Sosa reiterates his claim that *subcredal animal knowledge* requires just apt affirmation, and so ‘requires only guessing’ (2015, p. 151). Sosa’s notion of subcredal animal knowledge is not to be confused with what Goldman and Olsson (2009, §1) call ‘weak knowledge’, which needn’t be reliable, but only true. Goldman and Olsson, drawing from examples such as those offered by Hawthorne (2002), point out that it is sometimes felicitous to attribute knowledge to individuals who generate the right answer to a question, where the possession on this information contrasts saliently with *ignorance* (2009, pp. 19–20). On a further point, I note that the case of Simone is extremely rich and in many ways I think more complex than his other examples; the case is one I diagnose in detail in a separate work.

⁸ As Sosa points out, more confidence is something that can be gained through therapy, with no other change. ‘That would not give us the sort of knowledge we enjoy with the larger letters. Indeed such artificial increase of confidence can worsen the subject’s epistemic position’ (2015, p. 76). Even though the examinee does have more confidence in the case of the larger letters, it can’t thus be this confidence that explains the difference in epistemic standing.

Compare someone who gains not just confidence but also confirmation that he remains reliably right even when the letters shrink. This is perhaps the gain that raises him to a higher epistemic level. Now he might attain the knowledge requiring judgment, not just a guess. His knowledge would then comprise not only more confidence but also the proper meta-assurance that, even for those very small letters, his level of competence limits epistemic risk within proper bounds. Absent such additional confirmation, the vision-test subject lacks well-founded confidence on the second order that his first-order affirmations are more than sheer guesses. Even if by hypothesis his guesses are not right just by luck, he cannot be sure of that, not competently. By contrast, when the letters near the top are clearly and distinctly discernible, we do not just guess, which comports with our knowledge that we can identify those letters for what they are, that our affirmations at those rows would be apt (2015, p. 76).

In short, Sosa's reliable guesser lacks a certain kind of meta-assurance—viz., assurance that his level of (first-order) competence limits epistemic risk within proper bounds. The reliable guesser *ex hypothesi* can't tell whether his guesses would be, if right, right by luck.

But note that the second-order assurance the reliable guesser lacks (though which is *not* lacked when she affirms the larger letters on the top row) would not be secured simply by having more confidence *that* her first order affirmations are reliable enough.⁹ Such confidence could be incompetently, or unfoundedly, gained, as would be the case were the reliable guesser to believe, on the basis of a wish, that her first-order affirmations were reliable, even when they were reliable.¹⁰ Rather, what the reliable guesser lacks is knowledge that she can identify those letters for what they are, and *so that her affirmations at this row would be apt*. For all the reliable guesser can know, they would not be.

3 Performative defeaters

Sosa's reliable guesser in EYE EXAM is supposed to be on a kind of 'performative par' with two other characters from Chapter 3:

- (i) the basketball shooter (noted in Sect. 1) who aptly makes a shot from just barely above his threshold of reliability. The shooter, who succeeds in endeavour to make the basket, is unaware that he is above his threshold of reliability, and so too easily would have shot inaptly.
- (ii) a huntress who, despite drinking much wine, aptly shoots a scurrying rabbit in twilight fog. Sosa stipulates that the huntress's shot was apt—successful because adroit—although (like the reliable guesser and the basketball shooter) not reflectively competent, because the huntress is unaware that

⁹ This is different from having more confidence, of each guess, that it is correct.

¹⁰ See Sosa (2015, p. 75, fn. 14).

her shot in these conditions would be apt and has thus underestimated her prowess.

I think, though, that these three cases—the reliable guesser, the basketball shooter and the huntress—betray some subtle but ultimately important differences. The differences end up mattering, I'll suggest, for whether we should agree that the reliable guesser in EYE EXAM attains a form of animal knowledge, what Sosa calls 'sub-credal' animal knowledge, *even if* we grant Sosa's stipulation that the reliable guesser affirms the bottom-row letters aptly. And, the diagnosis for why this is so will lead us to a better account of guessing (or so I hope to show).

In order to draw out the differences between the three cases, it will be important to first note a key second-order difference between the case of the huntress and the basketball player. The basketball player is what we can call *second-order oblivious* to the threshold of his own reliability. As Sosa describes the case, the shooter:

shoots indiscriminately at distances too near his threshold of sufficient reliability [...] In that zone—barely above the threshold—he is still likely enough to succeed, even without knowing that he is. So, even in his ignorance he can still perform with animal aptness, but full aptness is now beyond him (2015, p. 70).

Being second-order oblivious about whether you are above or below your threshold for sufficient reliability is *different* from being *second-order mistaken* about your threshold for sufficient reliability, even though being either second-order oblivious or second-order mistaken about your threshold for sufficient reliability is enough to prevent your first-order endeavour from being reflectively competent and thus fully apt.

Sosa's shooter is second-order oblivious, although to find an individual who is second-order mistaken, we need only to consider again his huntress, Diana. As Sosa describes the case, Diana thinks there is little chance of success (in light of her having drunk much wine, and the twilight fog of the conditions), and in thinking this, 'underestimates her prowess' (2015, p. 68).

Unlike the basketball shooter, Diana *actually thinks that she's outside our threshold of reliability*, though she happens to be within it. The basketball shooter doesn't think this. The basketball shooter doesn't, like Diane does, underestimate (or for that matter overestimate) his prowess. He simply shoots aptly whilst *lacking any appreciation* that he is within his threshold for sufficient reliability.

The difference between the basketball shooter and Diana can be accounted for in terms of what we can call *performative defeat*.¹¹ In traditional epistemological discussions of defeaters,¹² defeaters are beliefs or doubts which count against (in the case of *rebutting* defeaters) the truth of a target belief, or (in the case of *undercutting* defeat) against the reliability of the target belief's formation.

Call a subject *S*'s belief that *p* a *performative defeater* if, with respect to some occasion of endeavour *E*, *S*'s belief that *p* (in the rebutting case) counts against *S*'s

¹¹ Cf., Carter and Navarro (2015) for a very different strand of performative defeasibility, one where the object of defeat at issue is specifically *knowledge-how*, on an anti-intellectualist construal.

¹² See, for example, Bergmann (2005), Lackey (2008), Pollock (1986), and Sudduth (2008).

attaining success in *E* and (in the undercutting case) calls into doubt the reliability of *S*'s competence in *E* by indicating that were *S* to succeed in *E*, the success would not be apt. Framed this way, it is of course possible for these defeaters to overlap. For example, suppose that the occasion of endeavour in question is that of sinking a simple two-foot putt. Suppose though that, through irrationality, a golfer believes that she has severe Ehlers-Danlos syndrome, an invisible but debilitating connective tissue disorder. The golfer's belief that she has Ehlers-Danlos is a rebutting performative defeater. It counts against her successfully attaining her aim (sinking the putt) on this occasion of endeavour. But it is also undercutting (in the sense described) because it indicates that *were* she to succeed, the success would not be apt—it would not be because of adroitness but because the ball happened to go in despite unreliable form.

Though rebutting and undercutting performative defeaters can also come apart. Consider that, as described, performative defeaters are psychological or mental state defeaters¹³: they defeat, as Lackey puts it, simply 'by being had', and not by being rational or true. Accordingly, there will be cases where the two kinds of defeat can come apart.

For example, suppose Sosa's skilled huntress, Diana, is told that a prankster has set up invisible forcefields around many, but not all, of the targets in the field. Suppose further the targets—with and without forcefields—are far away enough that only one with great skill, such as Diana, would be able to successfully hit one at such a range. By acquiring the belief that some of these targets are surrounded by invisible forcefields, the huntress acquires a performative rebutting defeater, one that counts against her succeeding on an occasion of endeavour to shoot and hit one of the targets in the field in these conditions. *However*, this belief arguably does not indicate that *were* the huntress's shot to be successful, that it would thereby not be apt. Such a successful shot would surely be creditable in large part to the huntress' impressive competence.¹⁴

Rebutting and undercutting performative defeaters can come apart in the other direction as well, where an agent has an undercutting performative defeater, but not a rebutting performative defeater. Singer Adele Adkins, for example, reports a dejected experience prior to singing on stage: though she does not have any doubts about how to sing the notes she is meant to sing—it is not as though she scrambles to double-check the lyrics or worries irrationally that she is suddenly too hoarse to sing them—she reports worrying, in a nebulous sense (perhaps irrationally), that *she* will bomb the performance. Adele, to stress, does not report having any beliefs, about any particular notes or songs, that she does not know how to sing them. She rather has doubts about herself—viz., that her performing a set of songs on stage will manifest her competence to sing them. Thus: an undercutting performative defeater in the absence of a rebutting performative defeater.

¹³ It is also possible to model rebutting and undercutting psychological defeaters as normative defeaters. Normative defeaters defeat not in virtue of being had but in virtue of being beliefs or doubts which one (epistemically) should have, regardless of whether one does in fact have them. For a helpful summary of the differences between these, see Lackey (2014).

¹⁴ For some established support for this line of thinking, see Haddock et al. (2010, Ch. 1–3), Pritchard (2012), Kallestrup and Pritchard (2012, 2013, 2014), and Carter and Pritchard (2015a, b).

With the foregoing picture of performative defeaters in mind, we can now diagnose in a more principled way the key difference between Sosa's basketball shooter, second-order oblivious, and the huntress, second-order mistaken, both of whom (like the reliable guesser in EYE EXAM) perform aptly at the first-order but, without second-order reflective competence, not fully aptly. Here is the difference: the second-order oblivious basketball shooter *does not have a performative defeater* (rebutting or undercutting),¹⁵ but the huntress does. The huntress has at least an undercutting performative defeater, given how Sosa has described the case, plausibly also a performative rebutting defeater.¹⁶ This is the case even though both are alike in that their respective performances are apt but not reflectively competent and so not fully apt.

Now, I think it matters for whether Sosa is right that his reliable guesser in EYE EXAM attains subcredal animal knowledge, whether the reliable guesser is more like the basketball shooter (who is second-order oblivious, with no performative defeater) or the huntress (who is second-order mistaken, with a performative defeater). In the next section, I want to argue that the reliable guesser is in fact like the huntress, not the basketball player, vis-a-vis performative defeat. This, it will be suggested, a consequence of the particular way that *guessing*, as a representational attitude distinct from ordinary belief, should be understood as having an alethic aim.

4 Guessing, the truth aim, and defeat

A point familiar from James (1897) is that the aim of possessing the truth is best understood as a mixture of *two competing aims*: truly representing and not misrepresenting.¹⁷ This would be relatively uninteresting if these aims were in no way in competition with one another.¹⁸

However, there is an important sense in which truly representing and not misrepresenting are *competing* aims. This point can be made in terms of risk: one can only fulfil the first aim by putting oneself at risk of not fulfilling the second. Weighting

¹⁵ At least, the huntress has one kind of performative defeater, an undercutting performative defeater, which the second-order oblivious basketball shooter does not have. This point holds independently of the separate issue of whether either has a *normative* defeater. For the present purposes, and to avoid what is I think unnecessary complication in light of the present point being made, I remain neutral as to whether we should regard the second-order oblivious basketball shooter (or Diana for that matter) as having a normative defeater.

¹⁶ It is the huntress's explicit *underestimation* of her prowess that constitutes for her an undercutting performative defeater, in the sense described. This is compatible with her also having a rebutting performative defeater, which she would have if she, further, positively believed (or at least, doubted) that she would successfully hit the rabbit. Sosa's discussion of the case indicates that Diana probably also entertains such doubts, though this isn't made explicit. What is of particular relevance to the comparison between Diana and the case of the basketball shooter is that the basketball shooter, in virtue of being second-order oblivious, does not have either kind of performative defeater, whereas Diana has at least an undercutting performative defeater.

¹⁷ As James (1897) puts it: 'There are two ways of looking at our duty in the matter of opinion [...] We must know the truth; and we must avoid error—these are our first and great commandments [...]' For more contemporary expressions of this idea, see Alston (1985), Foley (1987), David (2001), Fallis (2006). See also Carter et al. (2015, 2013).

¹⁸ Or—in Smith (2014) terminology—if these aims were *normatively coincident*—viz., if they were such that one could not aim at one without automatically aiming at the other (2014, p. 273).

the second aim, not misrepresenting, more would lead one to be more cautious so as to avoid possible misrepresentation. Correspondingly, weighting the first aim more would lead one to be bolder in order to possess more truths.

In Carter, Jarvis and Rubin (Forthcoming), we showed how different ways of mixing these two aims correspond to doxastic attitudes at *different places in the hierarchy of representational attitudes*. Here is a brief overview of this idea:

There are attitudes like being-sure, where the second aim is appropriately weighted more (so that misrepresentation is worse and failing to truly represent isn't so bad), and there are attitudes like suspecting, where the first aim is appropriately weighted more (so that failing to truly represent is worse and misrepresentation isn't so bad). But, *the aim of all of these doxastic attitudes is to possess the truth by having the attitude*. It's just the mixture of the twin aims—how important they should be in different situations—that varies.

Guessing is very much like suspecting¹⁹ in the following sense: both are attitudes that, by taking them up, we aim at truth, but in a sense that is weaker, or less cautious, than belief.²⁰ With states like suspecting and guessing, more error possibilities are tolerable in order take a chance on possessing the truth.²¹ I won't rehearse our argument for this picture here.²² However, to the extent that the model proposed is a plausible one—viz., that in taking up the attitude of suspecting, or guessing, we aim at the truth in a way that affords more weight to truly representing than to not-misrepresenting (and vice versa for the attitude of being-sure-that)—we have cause to wonder whether 'knowledge' (animal or otherwise) is really what we should call whatever achievements ensue from reliable guessing or suspecting.

In order to appreciate this point, contrast Sosa's reliable guesser with a slight variation on EYE EXAM, involving a reliable *believer*. Suppose the reliable believer and the reliable guesser's first-order visual-perceptual competence is identical. Both affirm the top letters more easily. The reliable guesser affirms by resorting to guessing at the bottom row, the reliable believer affirms by believing. *What is the difference between the two?*²³ Saying there is no difference is of course not a viable

¹⁹ Sosa, in Chapter 8, revisits the eye exam case, and in his discussion, he notes that suspecting and assuming are *kinds* of guesses, guesses without affirmation. It's not clear to me whether suspecting, on the model advanced in Carter et al. (2013), should be thought of a non-affirmative attempt at attaining the truth, or rather, as a way of affirming which differs from ordinary (belief-based) affirming in the way that the twin goals of representing truly and not-representing are weighted.

²⁰ On Sosa's model, there are two primary varieties of belief: functional beliefs and judgments. The discussion of belief, in this section and with reference to the model advanced in Carter et al. (2013), pertains to belief, per se.

²¹ As it was suggested, misrepresenting because one is deceived by an evil Cartesian demon is an example of an error possibility that is tolerable for belief, but not for the attitude of being absolutely certainty. On the model advanced in Carter et al. (2013, §4), we suggested that by tolerating all error possibilities, one is giving maximal weight to truly representing, whereas by tolerating no error possibilities, one is giving all the weight to not misrepresenting.

²² See here also Carter et al. (2015, §3).

²³ It is worth considering here a difference between *both* guessing, and believing, and *estimating*. The difference is, I think, a matter of whether the individual takes a chance on possessing the truth *directly* or *indirectly*. The guesser takes a chance on possessing the truth *directly* by affirming. By estimating, one

option²⁴. But, the most viable way to account for what such a difference would be will ultimately be in tension with Sosa's diagnosis of the reliable guesser's achievement as animal knowledge (subcredal or otherwise). I'll now explain this point.

Firstly, consider that the reliable *believer* is best understood as on a performative par with Sosa's basketball shooter. Recall that the basketball shooter, who makes a shot from just above his threshold for sufficient reliability, but unaware of this, was second-order oblivious, *not* second-order mistaken, and this was the more fine grained characterisation of why the shooter lacked reflective competence. So described, we said that Sosa's basketball shooter did *not* have a performative defeater, as the huntress surely did. The same is the case for a 'belief' variation on Sosa's reliable guesser. That is, if we hold fixed the first-order competence of Sosa's reliable guesser in affirming the bottom row letters—in which case the letters still appear somewhat blurry, even though they are affirmed correctly—but suppose that the attitude taken up is a *belief* rather than a guess, we are now in the same position as the basketball shooter: we have a first-order apt performance that is not reflectively competent because the performance is second-order oblivious.

But Sosa's reliable guesser is *not* second-order oblivious, like she would be were she taking up the attitude of belief. And here I think (i) the heirarchical model of representational attitudes as a function of mixing the competing aims of truly representing and not misrepresenting, along with the notion of (ii) performative defeat introduced in Sect. 1, can help us to diagnose just why.

Here it will be helpful, initially, to consider the classic case of the chicken-sexer²⁵. This is paradigmatic case of an individual with a reliable competence but who is not reflectively competent because second-order oblivious. The chicken sexer purports to *believe* of each chick, what its sex is—the chicken sexer would *not tell you she was guessing*. She is *believing*. This is so even though the chicken sexer cannot furnish you with any grounds for her belief. As Sosa (2015, p. 76) rightly suggests, the chicken sexer attains an apt belief, and so animal knowledge.

Now, let's run a quick variation on the chicken-sexer case. Suppose that the chicken-sexer took herself to resort to guessing, rather than believing. Suppose she *tells* you she was guessing. Does she attain animal knowledge, provided we assume the first-order reliable competence (at sexing chickens) is the same as it is in the original case? Provided she is sincere, the answer should be no, and the explanation for this has two

Footnote 23 continued

affirms in the endeavour to be *near enough right*, where near enough right is the proximal goal, and truth is the distal goal. There is some philosophical precedent for regarding estimating as corresponding to a distinctive *faculty*, different from belief-forming faculties. The medieval philosopher Ibn Sina (Avicenna) defended such a distinctive estimative faculty under the description of *wahm*. However, Avicenna's characterisation of *wahm* leaves one to suspect that he had in mind a faculty that was fundamentally emotive rather than cognitive. An example Avicenna frequently uses involves an estimative judgment of hostility, as one might experience when nearby a wolf. In some passages, Avicenna explicitly divorces *wahm* from what he calls the intellect. See Avicenna (1027 [1952]).

²⁴ Sosa refers to subcredal animal knowledge as 'below even the animal knowledge that requires belief' (2015, p. 76), expressing his view that subcredal animal knowledge is not a *kind* of doxastic (belief-based) animal knowledge, but animal knowledge that arises from taking up the lower or 'weaker' representational attitude of guessing.

²⁵ See, for example, Goldman (1975).

parts. The chicken sexer, in regarding herself as guessing, plausibly has acquired a psychological defeater that is lacked by the *believing* chicken-sexer. Indeed, *unless* one is guessing while regarding oneself as believing, one is, resorting to guessing, acquiring a psychological defeater—namely, *that* she is guessing—which undermines (unless defeated) one’s animal knowledge.

Why is this? Why is mere *awareness* that one’s aim at truth is a guess, itself enough to constitute a defeater, one that suffices to defeat animal knowledge? The answer is that regarding oneself as guessing is at the same time regarding oneself as aiming at truth in a way that tolerates more error possibilities than are plausibly tolerated by knowledge.²⁶ Guesses can be apt, no doubt. But when they are, the ensuing achievement is not reliable enough to qualify as knowledge, animal or otherwise, because of the type of bold attempt at truth which a guess is, one which weights representing truly much more than not-representing falsely.²⁷

The basketball-performative analogue to the chicken-sexer-cum *guesser* and also to Sosa’s own reliable guesser will be a player who regards herself as having attempted to score in a manner that tolerates greater-than-average error possibility—by *shooting-via-chucking* the ball—than in normal circumstances. *Even if, by stipulation, the basketball player’s chucking form is actually reliable enough*, the fact that the basketball player regards herself as chucking the ball—a method she appreciates to be less reliable than usual—rather than shooting normally, the shooter acquires a performative defeater which renders the shot on par with the chicken-sexer-cum-guesser—namely that they are aiming at attaining the truth by guessing, and so thereby by tolerating greater than normal error.

Sosa’s reliable guesser in EYE EXAM is on a performative par with the following three reflectively incompetent individuals: the chicken-sexer-cum-guesser and the basketball player who acquires a performative defeater in light of knowingly *chucking* the ball,²⁸ unknowingly doing so reliably, and Diana, who is second-order *mistaken* by underestimating her prowess from the range at which she is shooting. *None of these characters is second-order oblivious*; each is second-order mistaken in a way that generates a performative defeater. The basketball chucker because he regards

²⁶ A corollary to this point is that, by endeavouring to possess the truth by guessing, one endeavours to possess the truth in the absence of normal default trust (see here Sosa 2015, p. 81).

²⁷ The foregoing discussion is compatible with granting that beliefs and guesses can potentially not differ in their reliability. However, as an anonymous referee at *Synthese* has pointed out to me, in some cases, coming to believe that you are guessing could actually have the consequence of reducing one’s reliability in the relevant domain, perhaps substantially. This is because regarding oneself as taking on above-normal levels of risk could potentially generate for one a behavioural confirmation effect, where the individual over time increases the level of risk undertaken, rendering her doxastic performances (in the relevant domain) less truth-conducive, and below the threshold of reliability required for first-order aptness. Though, even when such a behavioural confirmation effect does not materialise, one’s simply regarding oneself on a given occasion as affirming in a way that involves above-normal levels of risk is itself sufficient to generate a defeater for the agent’s would-be knowledge.

²⁸ Note that this case is *not* meant to be the same as a case where one chucks the ball from half-court. The latter is best construed as analogous to a *wild* guess, of which I saw more later. Rather, the case referenced above is one where we are to imagine that the shooter, from barely within his threshold of sufficient reliability, attempts to make the basket while using a chucking style, a style regarded by the shooter to be not sufficiently reliable, even though it in fact is.

himself as issuing a shot from what he takes to be a less-than-reliable enough form, and the chicken-sexer-cum guesser and Sosa's reliable guesser because both have a belief which counts *against* their affirmations being reliable enough—namely that *they are aiming at attaining the truth by guessing*, and so by tolerating greater than normal error.

There is an important disanalogy then between these three reflectively incompetent individuals and another three reflectively incompetent individuals: the (ordinary) chicken sexer, and Sosa's basketball shooter (unaware that he is just above his threshold for reliability), and as his EYE EXAM examinee *would be* were the examinee to have trafficked in beliefs rather than guesses.

5 Some new parallels

To the extent that the previous section's discussion is right, then if there is such a thing as sub-credal animal knowledge, it will not be acquired in cases where an individual, no matter how reliable she is in the relevant domain of endeavour, takes herself to be guessing. The way that, when guessing, we aim at truth is a way which gives disproportionate weight to representing truly versus not misrepresenting, and as such is distinguished as an attitude which, just by taking up this attitude, we take a chance on attaining the truth by tolerating more error possibilities than usual, more than is plausibly tolerated by knowledge. As such, by *regarding* oneself as guessing—that is, when one guesses and is not mistaken that one is guessing—one thereby acquires a defeater which one would need to defeat in order to attain animal *knowledge*.²⁹ One way to defeat this defeater would be to, say, learn from the eye professional that one is extremely reliable, such that then one *no longer regards herself as guessing* when she affirms the letters on the lower row, so that now she *believes* what she affirms at the bottom row.

I want to conclude by doing three final things. Firstly, to make some remarks about unreliable achievements, secondly, to sharpen the discussion of weighting representing truly more so than not misrepresenting, in the case of guessing, and thirdly, to advance—from within Sosa's own programme—some performative parallels that I think map on to a range of different kinds of guesses.

²⁹ Cf., however, Hetherington (2001, pp. 155–56) for some dissent. On Hetherington's view—what he calls *gradualism*—some guesses can qualify as knowledge, though what he calls 'bad knowledge', knowledge of the lowest grade, which requires only true belief. Though he recognises that this is a minority position. Other rare defenders of the view that true belief, by lucky guess or otherwise, secures knowledge are von Kutschera (1982), Sartwell (1992) and Meno of Thessaly, who briefly advanced the view that knowledge is mere true opinion in the *Meno*. See also Foley (2012) for a more sophisticated defence of this view, on which knowledge is secured by true belief in conjunction with adequate information. For a survey of various views which have, with different qualifications, come close to the view that knowledge is true belief, see Martens (2006) overview of 'epistemological minimalism'. For explicit criticism of minimalist approaches, see Lycan (1994) and Kvanvig (2003, Ch. 1).

5.1 Unreliable achievements

Even if guesses cannot aspire to knowledge, animal or otherwise, this is not to say that guesses cannot be *apt*. In fact, Sosa's reliable guesser's beliefs are surely apt, in that their accuracy manifests a first-order reliable competence.

In Carter et al. (2015), we suggested that, just as there are varieties of representational attitudes that aim at truth differently (by corresponding, in the proper governance of these states, with different ways in which the twin goods of attaining truth and avoiding error are weighted), there are also *varieties of cognitive achievements*. An advantage of our model was an explanation for why knowledge, cum cognitive achievement, must ordinarily be safe, and this is because of the particular kind of attempt at success that *belief* is—namely, an attempt that places the premium it does on avoiding error.

Some achievements, including physical and artistic achievements, are of course highly *unreliable*, issued by means that ordinarily wouldn't be successful. Here's one blatant such example, suggested by Turri (Forthcoming):

BABY STEPS: Geno is eleven months old. He's daily been gaining confidence in his ability to balance on two feet. Today he stood holding onto the couch when an object in Dad's hand commanded his attention. Inspired, Geno let go and, to his parents' delight, took his first two steps in Dad's direction (Turri, Forthcoming, p. 14).

Likewise, consider the case of a potter:

POTTER: Ben produces his signature pottery, which is highly regarded, by methods that very often produce cracked and broken failures. In order to produce a new work, he may, in the typical instance, have to make several (or many) attempts. Ordinarily, his ways of making pottery is unsafe.³⁰

BABY STEPS and POTTER are examples of a phenomenon perhaps best captured with reference to the case of Red Sox batting champion Ted Williams: arguably the greatest hitter in the history of baseball, and who hit safely only 40% of the time. What BABY STEPS, POTTER and the case of Ted Williams suggest is just that some achievements are unreliable, in that they are issued by unsafe methods.

In the cognitive case, this is to be expected as well. When, for example, Sherlock Holmes makes an *educated* guess³¹—one that turns out to be right—this is a cognitive achievement, one creditable to his abductive prowess, though it is not knowledge, animal or otherwise. This is the case *even if* Sherlock's guessing turns out to be as reliable as Sosa's reliable guesser in EYE EXAM. Sherlock, like the reliable guesser, provided they appreciate that they are guessing, acquire a defeater—that, in taking a chance on attaining truth, they've by guessing tolerated more error possibilities than in normal circumstances, and more than is plausibly tolerated by knowledge.

³⁰ This example is adapted from Carter et al. (2015, p. 1603).

³¹ I'll discuss different kinds of guessing, including educated guessing, in more detail later in this section.

5.2 Guesses and weighting

Suppose a fair coin is flipped, and then covered up. What is your *credence* that the coin has landed heads? It should be .5. You of course shouldn't *believe* that the coin has landed heads. And this is as it should be; belief corresponds with a mixture of the weighting the aims of believing truly and not misrepresenting which places a premium on not misrepresenting. In short, this situation is riskier than what befits belief.³²

But do you *guess* that the coin is heads? Perhaps only if there is some practical pressure to affirm one way or the other. Nothing you know favours heads over tails. But you can nonetheless, in this circumstance, take a reckless chance on possessing the truth—and if it turns up heads, you were right.

Call the representational attitude—one that gives a maximal weight to representing truly and affords no weight to not misrepresenting—a *blind guess*.³³ The performative analogue, in basketball, would be taking a shot, with no appreciation of one's orientation with respect to the basket—viz., a shot in total darkness. This is an attempt at success which tolerates nearly all error possibilities.³⁴ It is fitting that blind guesses are sometimes referred to as *shots in the dark*.

But not all guesses are blind. Sosa's reliable guesser is not after all, *entirely insensitive to the aim of misrepresenting*, even if he takes himself to be guessing. Granted during the examination, as the letters get smaller, Sosa describes 'I start to lose confidence that I am getting the letters right'. But plausibly *some* confidence is retained, such that from the perspective of the guesser, some error possibilities are excluded. As Sosa describes the case, 'there are many cases where I am quite unsure as to whether it is an "E" or an "F", say, or a "P" rather than an "F", etc. (2015, p. 74).' This is a familiar situation: one where we affirm, due to some practical pressure to do so, while aware that some, though not all, relevant error possibilities remain live. Sosa's guesser is *not*, plausibly, unsure whether the letter before him is an "I" rather than an "F".

Guessing is a genus with different species. Guessing, as such, involves taking chance on attaining truth in a way that mixes the aims of truly representing and not misrepresenting *less cautiously* than *belief*, and in doing so tolerates comparatively more epistemic risk. Guesses can be *more or less* cautious, and I'll elaborate on this point in the next section by drawing some athletic-performative analogues.

It is worth highlighting that the foregoing characterisation of guessing is *compatible* with Sosa's account of the normativity of guessing—viz., as affirmation in the endeavour 'to get it right, without affirming in the endeavor to get it right *reliably enough*'—even

³² Plausibly, the appropriate attitude here is withholding, or suspending belief or judgment. For a recent critical discussion of what such an attitude involves, see Friedman (2013).

³³ Sosa himself discusses this kind of guessing under the description of a *sheer guess*, which is not what the reliable guesser is doing (see Sosa 2015, p. 151, fn. 14).

³⁴ The range of all error possibilities that's applicable depends in part on the relevant *whether* question. For example, you can take a blind guess and thus tolerate all error possibilities on a multiple choice test—choose A, B, C or D—by tolerating just three error possibilities. Blind guessing a given rational number, to be randomly generated by a machine, tolerates a potentially infinite number of error possibilities. In both cases, though, the blind guess as such is a chance on possessing truth that embraces maximal risk in the context of the guess.

though the foregoing is *incompatible* with the suggestion that reliable guessing aspires to *knowledge* (subcredal, animal or otherwise).³⁵

5.3 Performative analogies

Recall now some performative parallels that have already been drawn: a blind guess is akin to taking a shot at making a basket in conditions where one lacks any appreciation of one's orientation to the basket—viz., a shot in the dark. A more cautious attempt at gaining a truth is a *wild* guess, which involves affirming in the endeavour to get it right, but while tolerating *more* error possibilities than does Sosa's reliable guesser, when the reliable guesser affirms while unsure whether a given letter is an "E", an "F" or a "P". The wild guesser by contrast affirms (for example) while unsure, of a given letter, whether it is (in this case) *any somewhat wide letter*, thus excluding letters such as "I". If the guesser cannot even rule out the "I", and so affirms without being able to rule out *any* contextually relevant error possibilities,³⁶ the guess is not wild but blind. I suggested in Sect. 1 the basketball-performative analogue to a wild guess is a shot from half-court, under normal conditions; one can control one's shot choice so as to exclude shooting *away* from the basket, which cannot be ruled out when taking a shot in the dark, while nonetheless knowingly taking a chance on a basket that tolerates *most* error possibilities.

It's convenient to give guesses such names as *blind* and *wild* when they tolerate *all* or *most* error possibilities, by involving affirming in a way that affords either *no* or *barely any* weight to the aim of *not misrepresenting*.

But there are plenty of other varieties of guess which are less easily named—though a helpful way to think about them is by juxtaposing them with their practical performance analogues. Specifically, we can do this by considering cases where the performative defeaters in cognitive and non-cognitive cases line up.

Sosa's own *reliable* guesser was, I suggested, analogous not to the basketball shooter, second-order oblivious, who shoots from just above his threshold for reliability. That basketball shooter did not *have* a performative defeater. The shooter

³⁵ It is worth noting one potential line of reply here, to the effect that defeaters which defeat *credal* knowledge (either, animal or reflective) do not defeat *subcredal* animal knowledge. If this were the case, then, plausibly, a reliable guesser could be a candidate for subcredal animal knowledge. There are at least two problems which would face any attempt to defend such a line. Firstly, one might endorse this line because one thinks that subcredal animal knowledge is *indefeasible*—immune to being undermined by ordinary mechanisms of defeat. Without some principled reason to treat subcredal knowledge as special in this respect, this line does not look promising. A slightly weaker way to suggest that defeaters which defeat credal knowledge do not defeat subcredal knowledge would need to make some kind of argument from strength—to the effect that defeaters which suffice to defeat credal knowledge are not strong enough to defeat subcredal animal knowledge. This is weaker than the previous envisioned line because it is not committed to the indefeasibility of subcredal knowledge. However, on closer inspection, the weaker line looks similarly unprincipled. After all, it would be mysterious indeed why the title of 'knowledge' should be retained in the presence of undefeated defeaters (of the sort which defeat credal knowledge), even if some positive subcredal epistemic status were retained. Thanks to Modesto Gómez Alonso.

³⁶ In the context of an eye examination, the contextually relevant error possibilities are restricted to possible letters. It's common knowledge, in this context, that error possibilities do not include things such as numerals and non-letter shapes.

would, however, were he to have *chucked* (rather than shot normally) the ball in a way that was *unbeknownst to him* reliable enough. That's the analogue to reliable guessing.

We can distinguish Sosa's reliable guesser from a guesser in a slightly different situation, which is best appreciated by considering first an athletic-performative parallel. Suppose firstly that Sosa's basketball shooter is at a range *very slightly below* his threshold for sufficient reliability, but that he *knows* that this is the case, and takes (and makes) the shot anyway. This is similar to, but not quite like, the case of the individual who chucks the ball, unbeknownst to her, reliably. It's similar in that both regard themselves as tolerating above-normal risk (the former because she regards chucking to fall short of reliable enough form, the latter because she knows that from that distance, using normal form, she's slightly below the threshold for sufficient reliability). But the cases are different in that the shooter—shooting normally but knowingly slightly below his threshold for sufficient reliability—is (unlike the chucker who is unknowingly reliable) *actually* not sufficiently reliable, even though the shot is successful. So *reliable guessing* cannot be the right cognitive analogue performance.

The right analogue here is, I think, what we can call a (mere) *educated guess*. This will be when one, for example, hears from the reliable testimony of the optician that *the next row* is (based on a very recent calculation) very near, but just below, her threshold for sufficient reliability, and indeed the individual affirms correctly, but through a competence she knows is just barely not sufficiently reliable³⁷.

A final species of guess, different from those noted so far, that I think merits attention is analogous to a more complex—but not entirely uncommon—kind of performative situation. Consider the following:

3-POINTER: Steph, an excellent 3-point shooter, often reads his detailed shooting statistics to assess whether he is reliable enough from various places outside the 3-point line, to warrant taking open shots. Steph had recently read a statistical analysis of his performance from various ranges. The source he consulted, <http://www.fivethirtyeight.com>, uses expert Bayesian modeling, and Steph has no reason do doubt the source's reliability.³⁸ Steph is confident that if this source says he is above his threshold for sufficient reliability at a given range, then he is. Accordingly, on the basis of this source, Steph believes that he *is* reliable enough from 28 feet to warrant taking an open shot at 28 feet. However, Steph's coach also read the piece, and told Steph prior to today's game, that he is sure the number was 27, not 28. Steph thinks his coach is almost certainly wrong, and continues to believe that the number is 28—but nonetheless, in light of being unable (on the court) to dismiss this error possibility, Steph accepts that he does not *know* that he is reliable enough from 28 feet, despite his continuing to believe this. He takes an open shot from 28 feet and makes it.

This case is unlike any of the others so far. Firstly, like the reliable guesser, Steph *is* sufficiently reliable from 28 feet. (Suppose that Steph's coach was incorrect). But,

³⁷ This bears a close semblance to the everyday usage of 'educated guess' wherein one takes a chance on truth while believing she is tolerating above-normal risk in doing so, though below the level of above-normal risk one takes when chancing on truth via a wild guess.

³⁸ <http://fivethirtyeight.com/features/stephen-currys-bombs-are-too-good-to-be-true/>.

unlike the reliable guesser, Steph *believes* he is reliable enough from 28 feet, despite (in light of the on-court incorrect testimony from his coach) believing *that he does not know this*. Like the educated guesser, Steph has sought out reliable independent information about his reliability range. Unlike the educated guesser, Steph believes, before taking his shot, that he is from that distance reliable enough.

Does Steph have a performative defeater, in the sense articulated in Sect. 2? Yes. The on-court belief he acquires, that his coach insists the article noted 27 rather than 28 as his reliable-enough threshold, counts against the reliability of Steph's competence from 28 feet by indicating that were Steph to succeed in from 28 feet, the success would not be apt. Steph, it was stipulated, has doubts about the accuracy of this defeater, but—standing on the court—he does not regard himself as able to rule it out, to *defeat* the defeater, and so regards himself as not *knowing* he is reliable enough from that range, despite believing that he is.

What species of guess will be the cognitive parallel to 3-POINTER? In order to get this right, let's try to model this kind of performative defeater Steph has acquired as closely as we can, in the cognitive case:

PI: Elodin, a gambler, is attempting to memorise π to at least one digit past any of his friends, without making any errors (a single error nullifies the recitation, disqualifying him from winning the bet). Elodin has Cartesian certainty that the first three numbers are 3.14. Beyond that, his confidence ever so gradually wanes. Elodin's record, in practice, is 60 digits before failing. But, after practicing many times, he's concluded that he is reliable enough, in the context of the bet, only to attempt to recite π up 40 digits. The 40th digit of π is a "1". the 39th digit is "7."³⁹ Prior to the bet, a friend tells him that his calculations of his own reliability are very slightly off and that, on the basis of Elodin's practice records, he's in fact reliable enough only to attempt the first 39 digits, not the 40th. Elodin believes his friend is wrong, *but*, he realises he cannot conclusively rule this out. Elodin, believing—but believing that he does not know—that the 40th digit of π is a "1", does not stop at the 39th digit. He affirms, correctly, that the 40th digit is a "1".

Elodin affirms the "1" *while believing he does not know he is reliable enough* (given that knows he can't rule out that his friend is right, despite believing his friend is wrong) in the context of his affirmation. This affirmation is as such a species of guess: an attempt at truth whereby one takes a chance on gaining a truth in a way that tolerates more than usual error. But the risk Elodin knowingly undertakes here is marginal. He does not *believe* he is *not* reliable enough; he merely believes he does not know he is reliable enough, while believing he *is*.

Call this species of guess a *conjectural affirmation*. In the case of Elodin, the doubt he has, though it is appreciated as undefeated in the context of the affirmation, is compresent with his competent belief *that* he is reliable enough. Elodin, as well as Steph, after all, *believe* on good grounds that they are reliable enough, and affirm and shoot respectively while maintaining they do not know they are reliable enough.

³⁹ The first 40 digits of π are: 3.1415926535897932384626433832795028841971.

Conjectural affirmations, like other guesses, do not aspire to knowledge (animal or otherwise).⁴⁰

6 Concluding remarks

In the foregoing, I've shown why Sosa's view of guessing mistakenly identifies certain kinds of guesses as candidates for knowledge. I've argued for a different way of thinking about guessing, as a kind of attitude that by its very nature incurs a kind of defeater which is acquired when one takes a chance on truth by tolerating above-normal levels of risk. The rationale for the view advanced draws importantly from the Jamesian insight that the aim of possessing the truth is best understood as a mixture of *two competing aims*—truly representing and not misrepresenting. After situating my favoured approach to the attitude of guessing in the wider context of unreliable achievements, I concluded by distinguishing various species of guesses in connection with what are presented as their practical performative analogues.

I want to close by noting that I believe Sosa's overarching project in *Judgment and Agency* is a tremendous achievement in epistemology. It is, I think, his most compelling presentation of his distinctive brand of virtue epistemology to date. The aim here has not been to challenge Sosa's wider framework, as such, but to suggest how—from within this rich framework—a better account of guessing can be advanced, one which squares better with plausible insights about knowledge and epistemic defeat.⁴¹

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

⁴⁰ An interesting application of this line concerns the phenomenon of stereotype threat, which occurs when individuals who belong to negatively stereotyped groups regard themselves as likely to perform in accordance with the negative stereotype, and this tendency increases when one reflects on one's membership in a stigmatised group [see, for example, [Alfano \(2014\)](#)]. In extreme cases of stereotype threat, an individual from a group stereotyped as lacking in intelligence might, upon reflecting on her group membership, positively believe that she is cognitively unreliable at whatever task she is pursuing. If such an individual affirms nonetheless, her attempt at truth incurs a defeater and so she lacks knowledge (lest she can defeat the defeater, for instance, by coming to understand the mechanisms of stereotype threat). The perniciousness of stereotype threat becomes even clearer, though, when we consider less extreme cases, which are more in keeping with conjectural affirmations, such as in the case of Elodin. In such cases, suppose the marginalised individual, upon bringing to mind her marginalised group membership, simply comes to doubt her reliability (to some extent) while retaining her belief that she is reliable enough. In such a circumstance, even such an individual's belief is not (on the view I'm proposing) a candidate for knowledge for the same reason that Elodin's conjectural affirmation is not a candidate for knowledge. This is a frustrating but I think correct consequence of the view advanced here, one which highlights a special case of epistemic injustice. As [Alfano \(2014\)](#) has indicated, a way to combat this injustice is to give such marginalised individuals the capacity to defeat the kind of defeaters they acquire when subjected to stereotype threat, by learning about how stereotype threat works. Thanks to an anonymous referee at *Synthese* for suggesting I address this issue.

⁴¹ I am grateful to Modesto Gómez Alonso for helpful discussions and comments on previous versions of this paper. Thanks also to Emma C. Gordon, Jesper Kallestrup, Christoph Kelp, Joost Ziff, an audience at an Edinburgh University epistemology reading group, and two anonymous referees at *Synthese*.

References

- Alfano, M. (2014). Stereotype threat and intellectual virtue. In F. Owen & F. Abrol (Eds.), *Naturalizing virtue* (pp. 74–155). Cambridge: Cambridge University Press.
- Alston, W. P. (1985). Concepts of epistemic justification. *The Monist*, 68(1), 57–89.
- Avicenna. 1027 (1952). *Kitab Al-Najat (the book of salvation)* (F. Rahman, Trans.). Oxford: Oxford University Press.
- Bergmann, M. (2005). Defeaters and higher-level requirements. *Philosophical Quarterly*, 55, 419–436.
- Carter, J. A., & Navarro, J. (2015). The Defeasibility of knowledge-how. Unpublished Manuscript.
- Carter, J. A., & Pritchard, D. (2015a). Knowledge-how and epistemic luck. *Noûs*, 49(3), 440–453.
- Carter, J. A., & Pritchard, D. (2015b). Knowledge-how and cognitive achievement. *Philosophy and Phenomenological Research*, 91(1), 181–199.
- Carter, J. A., Jarvis, B., & Rubin, K. (2015). Varieties of cognitive achievement. *Philosophical Studies*, 172(6), 99–181.
- Carter, J. A., Jarvis, B., & Rubin, K. (2013). Belief Without Credence.? *Synthese*, 1–29. doi:10.1007/s11229-015-0846-6.
- David, M. (2001). Truth as the epistemic goal. In M. Steup (Ed.), *Knowledge, truth and duty: Essays on epistemic justification, responsibility and virtue* (pp. 151–169). Oxford: Oxford University Press.
- Fallis, D. (2006). Epistemic value theory and social epistemology. *Episteme*, 2(3), 177–188.
- Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.
- Foley, R. (2012). *When is true belief knowledge?*. Princeton: Princeton University Press.
- Friedman, J. (2013). Suspended judgment. *Philosophical Studies*, 162(2), 81–165.
- Goldman, A. (1975). Innate knowledge. In P. S. Stephen (Ed.), *Innate ideas* (pp. 20–111). Berkeley: University of California Press.
- Goldman, A., & Olsson, E. J. (2009). Reliabilism and the value of knowledge. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value* (pp. 19–41). Oxford: Oxford University Press.
- Haddock, A., Millar, A., & Duncan, P. (2010). *The nature and value of knowledge: Three investigations*. Oxford: Oxford University Press.
- Hawthorne, J. (2002). Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2), 517–523.
- Hetherington, S. (2001). *Good knowledge, bad knowledge: On two dogmas of epistemology*. Oxford: Clarendon Press.
- James, W. (1897). *The will to believe*. New York: Longmans, Green; Company.
- Kallestrup, J., & Pritchard, D. (2012). Robust virtue epistemology and epistemic anti-individualism. *Pac Philos Q*, 93(1), 84–103.
- Kallestrup, J., & Pritchard, D. (2013). Robust virtue epistemology and epistemic dependence. In H. Tim & P. S. David (Eds.), *Knowledge, virtue and action essays on putting epistemic virtues to work*. London: Routledge.
- Kallestrup, J., & Pritchard, D. (2014). Virtue epistemology and epistemic twin earth. *European Journal of Philosophy*, 22(3), 335–357.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- Lackey, J. (2008). *Learning from words: Testimony as a source of knowledge*. Oxford: Oxford University Press.
- Lackey, J. (2014). Socially extended knowledge. *Philosophical Issues*, 24(1), 282–298.
- Lycan, W. (1994). Sartwell's minimalist analysis of knowing. *Philosophical Studies*, 73, 1–3.
- Martens, D. (2006). Confidence in unwarranted knowledge. *Erkenntnis*, 65(2), 143–164.
- Pollock, J. (1986). *Contemporary theories of knowledge*. Savage, MD: Rowman & Littlefield Publishers.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3), 247–279.
- Sartwell, C. (1992). Why knowledge is merely true belief. *The Journal of Philosophy*, 89(4), 167–180.
- Smith, M. (2014). Knowledge, justification and normative coincidence. *Philosophy and Phenomenological Research*, 89(2), 95–273.
- Sosa, E. (2010). *Knowing full well*. Princeton: Princeton University Press.
- Sosa, E. (2015). *Judgment and agency*. Oxford: Oxford University Press.
- Sudduth, M. (2008). Defeaters in epistemology. *Internet encyclopedia of philosophy*. <http://www.iep.utm.edu/ep-defea/>.

Turri, J. Forthcoming. Knowledge as achievement, more or less. In M. A. Fernandez (Ed.) *Performance epistemology*. Oxford: Oxford University Press.

von Kutschera, F. (1982). *Grundfragen Der Erkenntnistheorie*. Berlin: Walter de Gruyter.