

# Can Performance Epistemology Explain Higher Epistemic Value?

**Abstract.** *Judgment and Agency* contains Sosa’s latest effort to explain how higher epistemic value of the sort missing from an unwitting clairvoyant’s beliefs might be a special case of performance normativity, with its superior value following from truisms about performance value. This paper argues that the new effort rests on mistaken assumptions about performance normativity. Once these mistaken assumptions are exposed, it becomes clear that higher epistemic value cannot be a mere special case of performance normativity, and its superiority cannot be guaranteed just by truisms about performance value. §§1-2 set the stage, clarifying the thesis and the relevant features of Sosa’s strategy, and explaining why the strategy requires the mistaken assumptions. §3 presents a dilemma for the new account of higher epistemic value. §4 deepens the case for one of the horns. §5 takes stock and draws some broader morals.

## 1 Introduction

According to performance epistemologists, all epistemic normativity is a special case of *performance normativity*.<sup>1</sup> Archery provides the favorite illustration of this exotic-sounding but familiar phenomenon. As Sosa (2007: 22) notes, we can assess an archer’s shot *qua* shot in several ways, asking whether it was

*accurate*—i.e., whether it hit the bullseye,

*skillful* or ‘*adroit*’—i.e., whether it manifested the archer’s target-hitting skill, and

*apt*—i.e., whether its very accuracy manifested the archer’s target-hitting skill.

Performance normativity is the genus of which such evaluative properties—the ‘AAA’ of accuracy, adroitness, and aptness—are species. These evaluative properties track the extent to which performances within a performance domain (e.g., archery, cooking, thinking...) are good *as* performances aiming at the fundamental objective(s) of that domain. As Sosa stresses on many occasions,<sup>2</sup> whether a performance is good *as* a performance

---

<sup>1</sup>‘Performance normativity’ is Sosa’s term. ‘Performance epistemology’ is a helpful new name from Vargas (2016) for what has longer been called ‘reliabilist virtue epistemology’.

<sup>2</sup>See, for example, Sosa (2007: 90), (2011: 59), (2015: 49), and (2017: 110).

aiming at some objective is independent of whether it is good *simpliciter* or all-things-considered. The domain-relative evaluations of performance normativity can hence be understood, as Ridge (2013) helpfully observes (and Sosa has confirmed),<sup>3</sup> as examples of *attributive evaluation* in Geach (1956)'s sense, in effect evaluating whether performances aiming at a target are *good shots* at that target. For Sosa, the fundamental cognitive target is truth.<sup>4</sup> Hence, cognition on his view fundamentally amounts to intellectual archery at the mark of truth.

Performance epistemology has greatly illuminated the nature of knowledge. Arguably the best account of knowledge around is Sosa (2007)'s account of knowledge as apt belief. Even those who worry that the account faces counterexamples concede the importance of aptness, accepting hybrid views that keep aptness as a central component.<sup>5</sup> Defenders of undiluted versions of performance epistemology have,<sup>6</sup> however, not obviously fared better than simple reliabilists in analyzing epistemic properties for which internalist intuitions are stronger than they are for knowledge. As Sosa recognizes, seemingly benighted subjects like Norman the clairvoyant, TrueTemp, and other epistemic subjects who lack good reasons (e.g., the superbindsighters of Block (1995)) can have beliefs that are adroit and indeed apt shots at the mark of truth in his sense. Yet even if one agrees with Sosa that these subjects have a kind of knowledge—viz., animal knowledge—one might think that there is something of great epistemic significance that they lack. To introduce some neutral terminology for adjudicating these issues, one might say their beliefs fail to exhibit any noteworthy *higher epistemic value* (HEV).<sup>7</sup>

---

<sup>3</sup>Ridge (2013: 198) suggests this interpretation, and says in note 19 that 'Sosa has confirmed this in conversation'. I have independently confirmed this interpretation in p.c. Sosa indicated that his only reservation about using Geach's notion of attributive value is that he wants to avoid making semantic claims or suggesting that his project is to understand ordinary language expressions of epistemic evaluations; the attributive/predicative distinction is, of course, a distinction concerning certain uses of 'good', 'better', etc.

<sup>4</sup>Performance epistemology can be combined with other views about the fundamental aim of belief, like Williamson (2000)'s view that the aim of belief is knowledge; for this version of performance epistemology, see Miracchi (2015a). Performance epistemologies with this conception of the fundamental aim of belief will, however, be less helpful if one is interested in giving a traditional analysis of knowledge.

<sup>5</sup>See, e.g., Pritchard (2012) and Kelp (2013).

<sup>6</sup>Some propose hybrid theories that incorporate themes of both reliabilist and responsibilist virtue epistemology at a basic level; Greco (2010), for example, gives an account of epistemic justification that appeals both to reliable cognitive abilities and to epistemic responsibility (without explaining it as a manifestation of the former). These theorists in effect concede that there is a dimension of epistemic evaluation that cannot be understood as a mere special case of performance normativity. The project of Sosa (2015) is, if successful, a major improvement, since it *derives* responsibilist themes from a unified performance epistemology.

<sup>7</sup>There are doubtless several different kinds of HEV that the beliefs of the benighted lack. My minimal claim is that there is some HEV that performance epistemology doesn't adequately explain.

*Judgment and Agency* contains Sosa's latest effort to explain how performance epistemology might accommodate HEV, venturing to reveal how (a) it could be a special case of performance normativity and (b) how its superiority over animal knowledge could follow from truisms about performance value. While Sosa's latest efforts have yielded crucial insights, I remain doubtful that HEV is merely a special case of performance normativity, and that its superiority can be explained merely by truisms about performance value. With that doubt in mind, my aim is to argue that

T: Pure performance epistemology lacks the resources to provide a satisfying account of HEV as a special case of its proprietary kind of normativity.

I will defend (T) in two stages. Firstly, I will argue that Sosa's account of HEV's superiority over animal statuses as a special case of a more general pattern in performance normativity rests on a mistake in the theory of performance normativity. Once this mistake is revealed, it becomes unclear how the superiority of HEV could be helpfully explained merely by truisms about performance value, or how HEV could be a special case of performance value. Secondly, I will argue that even if this assumption about how performance normativity works were true, performance epistemology would still fail to provide a satisfactory account of HEV. Along the way, I will address some objections, and I will conclude by taking stock and explaining why I think a conclusion as sweeping as (T) follows (though my main focus is on Sosa's truth-first performance epistemology, befitting this special issue).

## **2 Preliminaries: The Evolution of Sosa's Account of HEV and a Mistaken Principle**

Before proceeding to the arguments, I want to explain the principle about performance normativity that I will oppose and why Sosa's account requires this principle to be true if the account is to succeed as an attempt to illuminate HEV as a special case of a broader pattern of performance evaluation.

To bring out the principle, it will help to briefly rehearse the evolution of Sosa's account of HEV over the last ten years. In describing Sosa's project, I will use the stipulative expression 'high-grade knowledge' to cover the desirable kind of knowledge lacking in unwitting clairvoyants but present in mature ordinary perceivers of the same facts, but I set aside yet loftier epistemic achievements like understanding for the purposes of this paper. While it is not Sosa's own expression, I use this expression because it is help-

ful to understand Sosa’s project over the last several books as an attempt to adequately characterize the desirable kind of knowledge that lower animals lack. I do not use the expression ‘reflective knowledge’ because Sosa now reserves it for something he regards as subsidiary albeit important (see, e.g., Sosa (2015: 84) and (2017: 95)), though one might use it more broadly. This shift in usage makes it helpful to have a theoretically neutral term for the kind of epistemic status Sosa had initially called ‘reflective knowledge’ but now calls (as a part of the theory) ‘knowing full well’ or ‘judgmental knowledge’.

In Sosa (2007, 2009), he identified the core kind of high-grade knowledge with *reflective knowledge*. And in Sosa (2007: 32), he identified reflective knowledge with ‘apt belief aptly noted’. Hence, an initial picture of high-grade knowledge in Sosa’s post-2000 work is:

**HK1:** S has high-grade knowledge that p of the core kind iff (i) S aptly believes that p, where (ii) S also aptly believes that (i).

As Miracchi (2015b: 357-8) observes, Sosa (2007: 24; 2009: 135-6) gives two other non-equivalent glosses on reflective knowledge, according to which reflective knowledge is (a) defensibly apt belief and (b) apt belief that one’s source or basis of belief is reliable. Miracchi suggests that (a) is really ‘the core of Sosa’s idea of reflective knowledge’ (358). This suggestion does, I think, make better sense out of Sosa (2009)’s project and also restores continuity with his yet earlier (1991) conception of reflective knowledge, which put a strong emphasis on coherence.

Nonetheless, Sosa continues to mention apt belief aptly noted as a contrast to motivate his new picture, in which talk of the ‘reflective’ now becomes peripheral. As Sosa (2011: 9) notes, apt belief aptly noted falls short of doing what one might have wanted reflective knowledge to do. This is because one can have an apt belief that p and aptly note as much without one’s belief that p being *guided* or *informed* by the higher-order knowledge. As he puts the more general thought behind this worry, ‘a performance could be both apt and meta-apt while still falling short in that it is not in virtue of being meta-apt that it is apt’ (2011: 9). Hence, Sosa (2011) recommends switching from characterizing high-grade knowledge of the core kind as apt belief aptly noted (which remains ‘reflective knowledge’ in a narrow sense) to characterizing it as *fully apt belief*:

**HK2:** S has high-grade knowledge that p of the core kind iff S’s belief that p is *fully apt*—i.e., iff S’s belief that p is apt because meta-apt.<sup>8</sup>

Meta-aptness in turn requires ‘competent risk assessment’ (2011: 9), and the relevant

<sup>8</sup>See Sosa (2011: 7, 9, and 10) for varying glosses to this effect.

explanatory relation expressed by ‘because’ holds when such risk assessment ‘helps to guide the first-order belief that is apt’ (2011: 13).

In his most recent work, Sosa (2015, 2017) defends a similar account, and continues to emphasize full aptness as his ‘lodestar concept’ (2015: 65). But he has new ideas about the nature of guidance, belief, and aptness. Sosa now distinguishes two kinds of beliefs—*functional beliefs*, which include some passive seemings, and *judgmental beliefs* (2015: 66-7). *Judgment* that p, on his new account, is ‘affirmation in the endeavor to affirm with apt correctness’, and *judgmental beliefs* are dispositions to judge in this sense (2015: 66). High-grade knowledge of the central kind in this new framework is apt judgmental belief, which is itself essentially aptly guided:

**HK2\***: S has high-grade knowledge that p of the core kind iff S aptly judgmentally believes that p—i.e., iff S is disposed to aptly affirm that p in the endeavor to so affirm aptly.

It is noteworthy that on Sosa’s current account, judgment doesn’t aim merely at truth but also at aptness.<sup>9</sup> For this reason, judgment itself is inapt if the judger fails to achieve aptness in the endeavor to affirm aptly: ‘[E]ven while affirming aptly in the endeavor to affirm correctly, he might fail to judge aptly. While affirming in the endeavor to affirm aptly, he might fail to affirm aptly in that endeavor. In other words, his alethic affirmation, aimed at truth, might be apt without being fully apt, in which case his judgment would not be apt’ (2015: 77).

Why think that high-grade knowledge *so understood* necessarily outshines mere apt belief? In appreciating Sosa’s answer to this question, we must pay close attention to the story in the text. It can be tempting to think that because Sosa is appealing to special kinds of activities and states—i.e., judgment and judgmental belief—he is making a move precisely analogous to a move Watson (1975) made in the post-Frankfurt (1971) compatibilist literature. Watson rejected Frankfurt (1971)’s account of free will in terms of higher-order desire, noting that we can raise all the worries that undermine a simple first-order desire compatibilism at the second order. Watson’s remedy was to appeal to a distinctive attitude of valuing that is not reducible to second-order desire, and which does not inherit its freedom from anything. One might imagine a version of Sosa’s project that makes the very same move. It is natural to worry in parallel to Watson that if high-grade

---

<sup>9</sup>Sosa (2011: 61) also claimed that when belief is understood as a disposition to affirm, it has the aim of aptness as well as the aim of truth: ‘Knowledgeable belief aims at truth, and is accurate or correct if true. It has accordingly the induced aim of attaining that objective. Such belief aims therefore not just at accuracy (truth), but also at aptness (knowledge). A belief that attains both aims, that of truth and that of knowledge, seems for that reason better than one that attains merely the first.’

knowledge is animal knowledge of animal knowledge, it is unclear why it should outshine mere animal knowledge. One might imagine someone suggesting *a la* Watson that the superior epistemic value of apt judgment comes for free with this distinctive attitude.

But this analogy is only partly apt. It is, to be sure, clear that on Sosa's new account, the agency involved in judgment is not reducible to anything at the animal level. Accordingly, his account is not reductive in the sense of explaining the greater epistemic value of judgmental knowledge in terms of further epistemic value at the animal level. But it does not follow that Sosa takes the special epistemic value of apt judgment for granted without deriving it from some more general principle. Instead, he continues to tell a 'special case' story within the theory of performance normativity: apt judgment is better than merely apt belief as a special case of a more general pattern of performance-normative evaluation. As Sosa (2015: 85) puts it:

The account of a desirable level of human knowledge as knowing full well is in fact a special case of something much more general.

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). A first-order performance will aim to attain a certain basic objective: hitting a target, as it might be. This will induce the correlated aim to succeed *aptly* in that basic aim. Attaining the basic success aptly is better than attaining it inaptly.

Moreover, aptly attaining the aptness of one's success is also better than attaining this inaptly. That is a lesson drawn from our basketball player near her threshold of competence, and especially from the amazing shot that spans the full length of the court to score a goal. This success suffers by comparison with that of a shot aimed not just at scoring but at doing so competently enough, and aptly. A shot suffers if it falls short of *full* aptness.<sup>10</sup>

Such claims would not be explicable by the theory of performance normativity unless we accepted the following principle for performance normativity, where 'upgrade' is only intended to caption the added performance-theoretic value mentioned in the principle:

**The Upgrade Principle:** An apt performance is a better performance in the relevant attributive sense if it is guided by apt awareness on the second order that the first order performance would be apt (likely enough).

Mark well that this principle does not suggest that a state or activity could be an animal-level state or activity (e.g., a functional belief) while also exhibiting full aptness, though

---

<sup>10</sup>While Sosa uses the word 'awareness' in the second paragraph, he is not assuming that the guiding state must have non-trivial phenomenology, nor that it is at the same judgmental level as the affirmation it guides; the relevant awareness may simply be doxastic awareness, and indeed of a functional sort. Elsewhere Sosa clearly allows that the relevant guiding state could simply be animal knowledge: 'Performing with full aptness would normally require knowing (at least at the animal level) that one would then perform aptly' (2017: 78). And Sosa has confirmed that this is the correct interpretation of his use of 'awareness' in p.c.

Sosa sometimes talks this way.<sup>11</sup> The principle is meant to allow that in the epistemic case, the performance would be judgmental in virtue of being an *affirmation* guided in the way that is characteristic of the fully apt. Judgmental aptness is, however, better in virtue of a more general pattern, which this principle partly underwrites. Indeed, if the Upgrade Principle were false, there would be no case for Sosa's claim that the fully desirable status for performances in general is full aptness *as a special case of a general pattern of performance evaluation*, though it might be brutally true or true for other reasons.

A complementary principle would support Sosa's remarks about how a performance can 'suffer' in performance-normative terms if it fails to be fully apt. Discussing Barney in fake barn country, Sosa (2015: 87) writes:

In falling thus short of knowing full well, does Barney fall short of something epistemically significant? In assenting here, I am guided by a more general, intuitively plausible principle: namely, that *any* performance suffers if it is not fully apt.

Like the earlier claim about betterness, this claim about worseness would be plausible if we accepted the following principle for performance normativity:

**The Downgrade Principle:** A performance is a lesser performance in the relevant attributive sense if it is not guided by apt awareness on the second order that the first order performance would be apt (likely enough).

Putting these ideas together, I will assume that in order for Sosa's story about why fully apt belief is better than merely apt belief to be plausible as a special case of principles *within the theory of performance normativity*, one must accept the following claim:<sup>12</sup>

**The Upgrade/Downgrade Principle (UP):** An apt performance is a better performance in the relevant attributive sense if it is guided by apt awareness that one is in a position to perform aptly, and a lesser performance in the relevant attributive sense if it isn't.

For again, if UP were false, it wouldn't follow from performance epistemology and its attendant scheme of performance evaluation that high-grade knowledge is as such better than animal knowledge. That comparative evaluative conclusion might be brutally true or true for other reasons, but *not as a special case of a broader pattern of performance evaluation*, as Sosa claims. Hence, UP's falsity would undermine the strategy for explaining

---

<sup>11</sup>Consider Sosa (2015: 54; my bold): 'Functional, perceptual seemings are passive states that we cannot help entering. But *endorsement of them* remains volitional, agential. And such endorsement is *required for those functional states to ascend to the level of fully reflective, judgmental knowledge*, the level to which the Pyrrhonists aspire, in which they are followed by Descartes.' In reading a draft of the book, I urged Sosa not to make this claim, and he does later back off (see Sosa (2015: 94)).

<sup>12</sup>If it isn't clear from my wording, note that I don't deny that principles of this form might hold for other forms of normativity. Indeed, the point is that performance normativity seems different.

HEV as a special kind of performance normativity. For this reason, my strategy for defending (T) involves discrediting UP.

### 3 The Dilemma and the Core Problem

#### 3.1 The Dilemma and Core Problem in Outline

I turn first to argue that UP is either false or does not relevantly help to illuminate the superiority of high-grade knowledge over animal knowledge. The argument takes the form of the following dilemma, where some crucial terms will become clear shortly:

1. If the bearers of performance-normative status are understood narrowly to include only deeds and attempts, UP is false.
2. If the bearers of performance-normative status are understood more widely to include aimings and policies, UP cannot helpfully explain the superiority of high-grade knowledge over animal knowledge as a mere special case of a pattern in performance normativity.
3. The bearers of performance-normative status are either to be understood narrowly, as deeds and attempts, or more widely, to include aimings and policies.
4. So, UP is either false or cannot helpfully explain the superiority of high-grade knowledge over animal knowledge as a special case of a pattern in performance normativity.

I will defend (1) and (2) directly. In defense of (1), I will present a counterexample, consider what Sosa has to say about counterexamples of this sort, and argue that his line on these cases is not persuasive. In defense of (2), I will give one main argument, which will bring out some further support for (1).

I will not directly defend (3). It is, I take it, uncontroversial that performances could be understood more narrowly or more widely. What sorts of things should go under the 'narrow performance' and 'wide performance' headings is, however, less obvious.



Here I am simply following Sosa (2015: 124)'s taxonomy.<sup>13</sup> Under the narrow heading, I am including at a minimum the things that we describe with action words in ordinary language, where such actions include both *completed doings of things attributable to agents* ('deeds' is Sosa's term) and *doings with aims, including doings-in-progress, that may fail to achieve the aim* ('attempts' is Sosa's term). Under the wider heading, I am placing the phenomena that Sosa regards as sufficiently like performances in the narrow sense to qualify as proper bearers of performance-normative status. These include 'aimings' (Sosa's term) in the intuitive sense that includes intentions, which embody our answer to the question, 'What are you aiming to achieve?'

Under the wider heading, we also find *policies*. Their inclusion deserves some comment. When considering the kind of counterexample that concerns me, Sosa (2015: 124-8) initially responds by distinguishing between the deed, the attempt, and the aiming as the main examples of performance kinds. But in reading an earlier version of this paper, Sosa suggested that in responding to my concerns now, he would want to understand stored belief and knowledge not as deeds, attempts, or aimings, but rather as akin to policies, where he is here understanding policies as dispositions to perform in certain ways. He makes this suggestion at one point in the book, though only in a late chapter.<sup>14</sup>

Now, I confess that I find it hard to see how policies are sufficiently performance-like even to be proper targets of performance-normative evaluation. The same goes for some aimings (viz., intentions). But my main concern here is not this one, but rather that the extension of performance epistemology needed to encompass all intuitive attributive evaluation of policies and some aimings is unprincipled.

Here is a preview of the main worry. Policies and intentions can be evaluated in many ways. If some knowledge were a kind of policy, some of these modes of evaluation might be used to explain why some forms of knowledge are better than others. But

---

<sup>13</sup>For the terminology that follows, see Sosa (2015: 124), though note that he subsumes attempts under aimings (or at least treats them as having the normative profile of aimings):

1. Performances are of two sorts: (a) deeds, and (b) aimings. Kicking one's spouse in one's sleep is a deed, an attributable doing (unlike squashing a rabbit at the end of one's fall off a cliff), but is not an aiming, something one aimed to do.
2. Aimings are of two sorts: (a) functional or teleological (whether biologically, socially, or psychologically, whether by a whole animal or by a subsystem), and (b) intentional.
3. Moving to epistemology, consider alethic affirmations, aimed at truth. An intentional action can have more than one end, constituting more than one *attempt*. So, an attempt constitutively involves a particular aim. It is this distinctive aim that makes it subject to the AAA normative structure of such attempts, which can be accurate, adroit, and/or apt.

<sup>14</sup>Sosa (2015: 210).

performance epistemology attempts to reduce all epistemic evaluation to a specific kind of evaluation—namely, attributive performance evaluation of the type paradigmatically expressed by words like ‘skillful’. Accordingly, performance epistemology cannot easily help itself to other modes of evaluation, especially not ones it wants to explain. It must be shown how performance epistemology’s own attributive evaluative principles, or natural extensions of them, help to explain the relevant normative facts. And I do not see that there is a principled extension of performance epistemology that does this work.

This is the main concern behind premise (2), and it is the core problem I am raising for performance epistemology in this paper. I agree that there is a sense in which intentions that are responsive to subjective reasons are better as intentions. I also agree that there is a sense in which policies that are responsive to subjective reasons are better as policies. But I do not think that these facts support Sosa-style virtue epistemology. These facts instead reveal its main mistake, which is to assume that the attributive evaluation of states and processes under our rational control is skill evaluation (or some natural extension thereof).

Here it helpful to compare epistemology to the philosophy of practical reason. Kantians and Aristotelians agree that normative practical assessment is grounded in the constitutive principles of intentional agency. As Korsgaard (2009: Ch.8)’s work indicates, Kantians and Aristotelians also agree that what makes an action or intention good is determined by whether that action or intention is good *as* an action or good *as* intention. But these approaches are also importantly different. For Aristotelians, all practical evaluation is to be understood as a kind of skill evaluation, where robustly normative practical evaluation is skill evaluation relative to the constitutive aim of human action (*viz.*, flourishing). Constitutive norms for Aristotelians are hence properly described as *teloi*, and practical goodness is skillfulness in attaining the *teloi*. Kantians do not understand constitutive norms as *teloi*, and hence are not committed to thinking of the attributive evaluation of action and intention just as a kind of skill evaluation (or some natural extension thereof).

Sosa’s performance epistemology is analogous to the Aristotelian approach to practical normativity. Like the Aristotelian approach, it is best understood as one that delimits the attributive evaluation of cognitive states and activities fundamentally to a kind of skill evaluation. I take my points to suggest that the attributive evaluation of cognitive states and activities goes beyond skill evaluation. I also take this claim to be plausible on its face. But it is not a claim that can be reconciled with performance epistemology unless performance epistemology is understood so widely as to include epistemic ana-

logues of Kantianism. That is not a helpful way of carving up the dialectical space. More importantly, that wider use of ‘performance epistemology’ would offer no escape route for Sosa, who is not a Kantian.

### 3.2 Defense of (1)

To see why we should accept (1), I will discuss a counterexample to UP that I raised in reading a draft of *Judgment and Agency*, a version of which appears in the book’s fifth chapter.<sup>15</sup> Here is the counterexample:

SKATEBOARDER: Dave was a pro skateboarder in his teens, but hasn’t been on a skateboard for fifteen years. Some old friends in town offer him a bet: they’ll give him \$50 if he can still do a certain trick as well as he did in his heyday, but he has to buy them all dinner at the most expensive restaurant in town if he can’t. Dave ought to worry that he’s forgotten how to do this trick—not because he’s in bad shape, but because he hasn’t been on a skateboard for years and it is beyond his budget to buy everyone dinner. But with brazen, unjustified confidence that he will succeed, he decides to try. As it turns out, he hasn’t lost his skill, and does the trick perfectly.

SKATEBOARDER looks like a counterexample to UP: the inaptness of the skateboarder’s belief that he can succeed does not seem to affect the attributive quality of his performance: he performs the trick perfectly, after all. Perhaps the skateboarder’s decision to perform is poor, relative to some appropriate scheme of evaluation. But the performance of the trick seems flawless in the relevant, attributive sense, nor would it have been a better performance of the trick if the relevant belief were apt.

These facts present the following problem. If the beliefs Sosa wants to classify as bearers of HEV were like tricks, they wouldn’t get any better or worse *qua* performances in virtue of being guided by apt or inapt beliefs or awareness of some loftier kind. Of course, one might insist that the quality of the believer’s decision to believe *just does* influence the evaluation of the belief. But this is beside the point: the point is that *if we take the driving analogies with performances seriously and focus on attributive performance evaluation, this is not what we should say*. We should rather conclude that such insistence involves the same error of which Sosa accuses some other virtue epistemologists: it assumes that adding goodies to the *auxiliaries* of a belief adds further goodies to the belief.<sup>16</sup> The analogy with other performances and their attributive evaluation suggests otherwise.

---

<sup>15</sup>See Sosa (2015: 125), where he attributes the example to me in note 15.

<sup>16</sup>See the critique of Baehr (2011) in Ch. 2.

Sosa (2015) responds in part on the basis of the observation that performances can be viewed in two ways—as deeds or as aimings, where among aimings he includes attempts.<sup>17</sup> The evaluation of aimings, Sosa claims, is plausibly sensitive to the aptness of the aimer’s guiding belief.<sup>18</sup> So, perhaps one can hold that if performances are aimings, UP is true, and the strategy for explaining HEV is vindicated.

Now, it is unclear to me that the performance-theoretic evaluation of *attempts* is so affected, and hence that the performance-theoretic evaluation of aimings in Sosa’s broad sense is also so affected. Consider the skateboarder again: his attempt was excellent in the attributive sense—a terrific success—despite his inapt guiding belief. So if aimings are understood in the broad sense that Sosa identifies, it is not clear that UP turns out true, or the strategy for explaining HEV vindicated.

I will consider what happens when we view performances as aimings in a more natural, narrow sense, which will turn us to the second horn of the dilemma. Firstly, however, I will undermine another move that Sosa uses to swing intuition in his favor. Toward the end of his discussion of this concern, Sosa brings up a different kind of example, which one might think suggests that there *must* be a deed-like kind of performance for which UP holds:

In a match, for example, the tennis player will always have various options when the ball comes at him across the net. Each option will have a risk/reward profile as the player aims for his relevant objectives, such as that of winning the point.

The good player will always be mindful of that profile, as he chooses stroke after stroke. So, his objective is not *just* attaining the desired immediate effect of the shot he settles on speedily. . . . No, that shot is in a deeper sense assessable with respect to whether it was a well-selected shot. If the opponent was far enough back from the net and the player was close enough. . . it may be blindingly obvious that he should hit an angled passing shot instead.<sup>19</sup>

My worry about this response is that the test case at the end is of the wrong sort. What we need is a case where the deed manifests first-order skill, but the doer has misleading evidence about her reliability. About the given intuition pump, we can simply say that the shot was flawed because although it succeeded, the player was in her position lacking reliable enough ability to be expected—objectively—to succeed. An epistemic analogue of *this* case is a case where a thinker who is reliable with respect to medium-difficulty questions chooses a maximum-difficulty question and answers it with the same confidence

---

<sup>17</sup>See Sosa (2015: 124–8).

<sup>18</sup>See Sosa (2015: 126).

<sup>19</sup>Sosa (2015: 127).

with which he answers medium-difficulty questions. While this thinker has skill with respect to medium-difficulty questions, his answer to the maximum difficulty question is *not competent at the first order*, since he *isn't sufficiently reliable* with respect to such questions. We get no support for UP from such cases.

If we reflect on the relevant test cases in the practical sphere and train our eyes on performance normativity, we find no evidence for UP. SKATEBOARDER was already such a case. We can imagine tennis examples if we like. Imagine a player who hasn't lost her tennis know-how, but has been given misleading evidence to think she's lost it, that she's in the wrong conditions for its competent exercise, or whatever. Suppose she brazenly decides to play anyway, thinking she'll win. Does the shoddiness of her guiding belief have any influence on our intuitions about the quality of her shots? When I consider this question with the relevant, attributive sense of 'quality' in mind, my intuition is that the answer is 'No'.

### 3.3 Defense of (2)

I turn to the other horn of the dilemma.

Suppose that instead of viewing the relevant bearers of HEV as performances in the narrow sense, we view them as akin to certain agential states of the performer. There are two subtly different versions of this proposal that one can extract from Sosa (2015). According to one version, the relevant bearers of HEV are *aimings*.<sup>20</sup> Now, it is worth noting that Sosa uses the term 'aiming' in two ways. On a narrower use, 'aiming' includes just certain agential states of the performer. We find this use of 'aiming' in sentences such as the following: 'Intentional action involves aimings that are full-fledged intentions' (2015: 9) and 'When one judges, one affirms with the aim to get it right aptly by so affirming (where aiming is more than just wishing or hoping)' (2015: 82). Here the paradigm example of an aiming would be an intention.

Sosa also often uses 'aiming' more broadly to include actions understood thickly, as partly constituted by aimings in the narrow sense. We find this wider use in sentences like these: 'We find unity across action, perception, and knowledge. All three are constituted by aimings, by performances with a constitutive aim' (2015: 25) and 'Performances are of two sorts: (a) deeds and (b) aimings' (2015: 124). The argument I will shortly develop extends to a view on which HEV resides in aimings understood as literal performances, as we will see in the next subsection. But I am first considering its application to aimings

---

<sup>20</sup>See Sosa (2015: 124-5) for an emphasis on aimings.

in the narrower, more intuitive sense that centrally includes intendings.

A slightly different version of this proposal appears most clearly in the book's ninth chapter, and Sosa has emphasized in conversation that on reflection he would prefer to use this proposal to answer my objection. According to this proposal, the bearers of the relevant HEV are *policies*. Consider, for example, this passage:

[A] disposition seated in the will can amount to a freely chosen or sustained policy. Take the policy of signaling your turns as you drive. This can lead to a choice, conscious or subconscious, to signal on a certain occasion as you approach a turn. The particular free choice then implements the policy, the freely chosen general commitment to signal your turns. You are disposed to signal as you approach your turns, and your particular chosen signalings are guided by that policy, and manifest that freely sustained disposition.

That same idea is applicable to the sorts of beliefs now in our focus. These are supposedly dispositions to affirm freely, and such a disposition can itself reside in the will; it can itself amount to a freely adopted evidential policy. These beliefs thus stand in contrast to animal beliefs acquired through the unfree "automatic" proper functioning of our cognitive mechanisms.<sup>21</sup>

On this proposal, judgmental belief and knowledge are constituted by such policies.

One might think that both of these proposals would straightforwardly avoid my objection. After all, it does seem to be a fact about the evaluation of intentions and policies that these are sensitive to the quality of the subject's risk assessment. An intention is worse as an intention if it is unresponsive to apparent risk (even if merely apparent), as is a policy as a policy.

But this response is unpersuasive for several reasons. Firstly, the force of the counterexamples doesn't turn on whether the bearer of HEV is *identified* by the performance epistemologist with a performance. It turns only on whether the bearer has what one might call the *attributive normative profile* of a performance—i.e., whether the bearer is something whose attributive evaluation works like that of performances.

Someone who holds that all epistemic normativity reduces to performance normativity must take this view. For performance normativity just is the kind of normativity distinctively exhibited by evaluanda with the attributive normative profiles of performances. To deny that epistemic evaluanda have the attributive normative profile of performances would hence be to abandon performance epistemology, or at least its central explanatory ambition of understanding epistemic normativity in terms of performance normativity. Yet evaluanda for which UP rings true aren't ones with the attributive normative profile of performances. That is the point of examples like SKATEBOARDER.

---

<sup>21</sup>Sosa (2015: 210).

Perhaps, one might say, policies lack the attributive normative profile of performances. Still, one might insist that to accommodate HEV, the performance epistemologist's project can be extended to include understanding some epistemic evaluation as the intuitive attributive evaluation of policy-like and intention-like items.

But it is not clear how to extend the project in this way. Any simple extension risks a proliferation of strange bedfellows. The intuitive attributive normative profile of policies is, after all, a hotly contested matter. Some contestants would obviously sharply diverge from Sosa. A comparison with the practical case makes the point stark. Kantian constructivists like Korsgaard (2008, 2009) agree with Aristotelian virtue ethicists that all normativity reduces to attributive normativity (or, in other jargon, to constitutive normativity).<sup>22</sup> They would insist that the attributive evaluation of practical policies is substantially a function of whether these policies are in accordance with the Categorical Imperative, rather than (e.g.) with whether they contribute to flourishing. By analogy, one can imagine an internalist epistemologist—an Epistemic Kantian, say—insisting that the constitutive norms with respect to which doxastic policies are to be evaluated will fundamentally invoke something not connected to truth as the archer's aiming is connected to hitting the bullseye.

Accordingly, if performance epistemology of Sosa's kind is broadened to include any approach on which epistemic evaluation reduces to the attributive evaluation of all broadly agential targets (including intentions and policies), then competing views that are deeply unlike Sosa's will count as versions of Sosa's performance epistemology. This result is absurd. Hence, addressing the present worry requires more than just insisting that the attributive evaluation to which Sosa wants to reduce epistemic evaluation includes the evaluation not only of performances but of policies. We need a more specific characterization of the relevant kind of attributive evaluation.

In the case of performances, it is easy to get a grip on the relevant kind of attributive evaluation, and to distinguish Aristotelian from Kantian views. We have dedicated language for it: the language of skill. The relevant attributive evaluations are evaluations of performances as more or less skillful. If skill evaluation extended to states such as policies, we might avoid welcoming strange bedfellows into the fold of performance epistemology by restricting focus to such evaluations.

It is, however, unclear that skill evaluation helpfully extends to policies and inten-

---

<sup>22</sup>I take constitutive and attributive normativity to be cognate, since the standards of evaluation relevant to whether something is a good F in the attributive sense are standards given by the *nature* or *constitution* of Fs, in a broad sense that includes their *essential function*.

tions. Googling ‘skillful intention’ and ‘skillful policy’ yields precious few results.<sup>23</sup> There is perhaps a derivative way in which a policy or intention can be more or less skillful, in virtue of being one whose proper execution would be a skillful action. Skill evaluations of plans work in this way, as far as I can intuit. A plan would intuitively be inept iff the perfect carrying-out of the plan would be inept. If the carrying out of a plan were of necessity perfectly skillful, the plan would be too.<sup>24</sup> Plans, policies, and intentions seem of a piece. If so, the skill evaluation of policies and intentions will track the skill evaluation of their perfect executions. UP’s failure to hold for one will imply its failure to hold for the other. To the extent, then, that I understand how policies and intentions can be evaluated for skillfulness, I don’t see that the results benefit the performance epistemologist.<sup>25</sup>

Hence, I conclude that it is not helpful for the performance epistemologist to identify the bearers of HEV with agential states of the performer such as intentions or policies. The attributive evaluation of intentions and policies is sensitive to quality of risk assessment, but not in a way that the performance epistemologist is entitled to exploit. For the *skill* evaluation of these states is not relevantly sensitive to quality of risk assessment, and the attributive evaluative remainder beyond skill is not native to performance epistemology, but rather to a constitutivist epistemology more analogous to Kantian ethics.

### 3.4 Added Support for (1): The Trouble with Thick Deeds

The reasoning from the preceding subsection also creates a problem for a version of the view that knowledge is literally akin to an action, which Sosa has affirmed particularly loudly in very recent work.<sup>26</sup> One might have complained that in discussing deeds and attempts in the skateboarder case, the intuitions I was eliciting traded on thinking about these deeds and attempts in abstraction from the intention with which they are executed.

---

<sup>23</sup>Most of the 690 results for the former appear in New Age literature, and many of the 2200 results for the latter are irrelevant ones in which ‘skillful’ is in fact modifying a later noun (e.g., ‘skillful policy *advocates*’ and ‘skillful policy *researchers*’).

<sup>24</sup>Mark well that I am not claiming here that good plans cannot be ineptly executed. I am claiming only that a plan is only as skillful as its perfect execution would be skillful.

<sup>25</sup>To be sure, plan *formation* can be evaluated for skill. But knowledge is not akin to plan formation. Even if it literally were a performance, it would resemble a plan’s execution, not its formation. And the skateboarder’s execution of the plan to perform *that trick*—the one the skateboarder performed and planned to perform (admittedly without appreciating its fineness)—manifests perfect skill, not upgraded by prior skillful plan formation. In any case, the invocation of policies is intended to accommodate the worry found in Chrisman (2012) that knowledge is a state that can be merely stored, as Sosa stressed to me in raising a worry in the ballpark of Chrisman’s.

<sup>26</sup>See (2017: 71): ‘Judgment and knowledge itself are forms of intentional action.’



Yet one might think that there is a more fine-grained notion of a deed or attempt according to which deeds and attempts are partly individuated by the intentions with which they are performed.<sup>27</sup> Call such deeds/attempts *thick deeds/attempts*. One might insist that the attributive evaluation of thick deeds and attempts is no less sensitive to risk-assessment than the attributive evaluation of its constitutive intention.

But again, while it is true that the attributive evaluation of thick deeds and attempts is no less sensitive to risk-assessment than its constitutive intention, this fact doesn't bespeak victory for the performance epistemologist. For the performance epistemologist is seeking to explain HEV in terms of a more specific kind of attributive evaluation—namely, skill evaluation. And the skateboarder's *intentional doing* of the trick—his thick deed—isn't intuitively less skillful in virtue of the unwise planning that went into it. It is a worse deed, but not a less skillful deed. One could stipulatively use 'skillful deed' to mean 'attributively good deed'. But in doing so, one would be papering over the fact that one is importing a kind of evaluation foreign to reliabilist virtue epistemology, and at home only in an at least partly Kantian epistemology.

Instead of stipulatively using 'skillful deed' to mean 'attributively good deed', one might instead insist that we can appeal to a richer notion of skill, or a skill to do something richer. Indeed, one might reasonably insist that some skills clearly exclude negligence with respect to one's subjective reasons. Surely, for example, a babysitter is not skillful if he ignores evidence, even if it is misleading, that the child being watched is about to be harmed and does nothing.

But this response is unconvincing. One problem is that the underlying reasoning is invalid. The mere fact that there are activities such that their skillful exercise necessarily excludes negligence doesn't imply that HEV is analyzable in purely performance-normative terms. After all, just make  $\phi$ -ing some activity that has non-negligence or sensitivity to subjective reasons as part of its constitutive aim. The existence of such activities and the corresponding skills doesn't relevantly ensure that HEV is analyzable in purely performance-theoretic terms. And it is equally clear that there are other activities  $\psi$  such that skillful  $\psi$ -ing is compatible with negligence or irrationality, as SKATEBOARDER already suggested. That fact makes it hard to see how performance-theoretic notions and non-normative notions alone could explain HEV, as Sosa (2015) reckons.

Even if all cognitive attitudes and activities were such that skillful  $\phi$ -ing were inconsistent with negligence or irrationality, the performance-theoretic analysis would remain

---

<sup>27</sup>Compare Korsgaard (2009: 8-14)'s distinction between acts and actions, where acts are behaviors and actions are acts-performed-for-a-particular-reason.

unsatisfying, barring a story about why the epistemic domain is special. For the practical domain is not analogous. Not all activities are such that their skillful performance is inconsistent with negligence or irrationality, even relative to their aims. This fact alone limits the explanatory power of performance epistemology, since there is not a suitably close parallel between epistemic evaluations and the performance evaluations of activities like archery. Given that competence in Sosa's sense doesn't exclude negligence, it is hard to see how evaluations of negligence and non-negligence, as well as rationality and irrationality, are supposed to be grounded in skill evaluation.

### 3.5 The Core Problem

I take it that Sosa-style performance epistemology is not an epistemic analogue of Kantian constitutivism, nor a hybrid theory that incorporates constitutive normativity in the Kantian sense alongside skill-based normativity. Otherwise it has undergone a much more radical transformation than is advertised in the book. The problem is that the attributive evaluative principles to which the performance epistemologist would need to appeal to get around my counterexamples are principles with a Kantian inflection. Intentions and policies are indeed better as such if they reflect not merely skill but sensitivity to subjective reasons, including misleading ones. Thick deeds and attempts are also better as such if they reflect not merely skill but sensitivity to subjective reasons. But this fact is a problem for Sosa-style performance epistemology, not a basis for explaining HEV on Sosa-style performance epistemology's own terms. For performance normativity of the kind that interests Sosa is the more circumscribed normativity of skill evaluation, which the skateboarder case suggests is unsuitable for explaining the HEV of epistemic deeds.

## 4 The Unhelpfulness of UP, Continued

I turn to deepen the case for the unhelpfulness of the Upgrade Strategy, arguing that even if some upgrade principle held for the kind of performance evaluation Sosa invokes, HEV would remain insufficiently illuminated.

To see what I have in mind, let's confine our attention to aimings, and grant for the sake of argument that the following claim about aimings is true when 'apt' is used in Sosa's sense:

**UP 2:** An aiming is more apt, or attains a fuller kind of aptness, if it is guided by apt awareness with suitable content, and less apt if not so guided.

If we understand judgments as aimings at the mark of truth, we can then derive the following:

**UP 3:** A judgment is more apt, or attains a fuller kind of aptness, if it is guided by apt awareness with suitable content, and less apt if not so guided.

Even if these claims captured truths about the kind of performance evaluation that interests Sosa, they might not illuminate HEV. Perhaps aimings are *apter* or *better in performance-theoretic ways* if guided by relevant apt states. Still, they might not automatically attain the HEV we want. Even if we stipulate that they are apter in the proprietary sense, it may remain unclear that they are epistemically better. Perhaps increases in Sosa-aptness don't guarantee increases in real epistemic value.

#### 4.1 Clairvoyants and the New View

We can see why one might draw that conclusion by reflecting on unwittingly reliable clairvoyants like Bonjour (1980)'s Norman. Now, Sosa (2015)'s framework does have some heartening new implications about clairvoyants. One thing that seems problematic about Norman is that he is *guessing*. Sosa can fairly say that his new account does not accord HEV to subjects who are guessing. For high-grade knowledge on the new account requires apt judgment, and apt judgments aren't just guesses. So, if the complete diagnosis of why unwittingly reliable clairvoyants' beliefs fail to attain HEV were that they are just guessing, the new framework would do the trick.

But this diagnosis is incomplete. Consider a newly minted clairvoyant—Normann—who affirms in the endeavor to affirm aptly that he has clairvoyant powers and who can detect which first-order affirmations of his manifest them via a higher-order skill at hitting the mark of truth regarding his abilities, though he has no reason for thinking that he has *this* skill. Even if Normann has no reason for affirming that he has clairvoyant powers, he can coherently do so. Normann does more than guess. Yet it is natural to think his affirmation doesn't hereby attain the kind of HEV we want. The mere fact that he is making a judgment, which is itself guided by an apt higher-order judgment and not merely a seeming, doesn't affect my intuition here.

Notice that I am not here assuming that judgmental knowledge is animal knowledge guiding animal knowledge. This is why I explicitly talked about affirmation at both levels. The point is rather that judgments as well as beliefs can be baseless while still manifesting perfect skill to hit the mark of truth aptly, and adding guidance by a baseless higher-order judgment doesn't seem to remove the blindness of that first-order affirma-

tion. The mere fact that judgment involves a qualitatively different kind of epistemic agency doesn't by itself make a difference to the thought that there is something important missing—namely, reasons. In this way, the discussion here is importantly dissimilar to the Watson/Frankfurt exchange in the compatibilist literature on free will. Of course, one might insist that one cannot skillfully judge that *p* in the absence of good reasons for judging that *p*. But this is not Sosa's own tack, as we will see shortly: while he may sound favorable to the idea at one point,<sup>28</sup> he denies that competent judgment requires reasons in considering objections.<sup>29</sup>

As it happens, it is not even clear in Sosa's most recent work that he requires that the apt state that guides one's apt affirmation in the case of full aptness itself be a state at the same level of epistemic agency: 'Performing with *full* aptness would normally require knowing (*at least at the animal level*) that one *would* then perform aptly.'<sup>30</sup> If so, the only difference that Sosa's new picture makes to the clairvoyance example is that it rules out versions in which the clairvoyant is guessing in the sense of affirming not in the endeavor to affirm aptly but only in the endeavor to affirm correctly. That is not difference enough.<sup>31</sup>

There is another way to reach our conclusion. We should distinguish subjective evidence-relative risk—i.e., the kind associated with probability relative to one's subjective evidence—and objective risk—i.e., the kind associated with objective probability. Suppose there is no objective risk one is in the wrong situation or shape to perform reliably enough, and that one has great first-order skill. Suppose moreover that one has a higher-order knack that makes one sensitive to this objective risk, inducing one to believe one's circumstances and shape are right for apt success. Compatibly with these suppositions, there might be misleading evidence indicating that one's circumstances and shape are off which one ignores, so that one fails to properly heed the subjective evidence-relative risk in judging. Accordingly, one can satisfy the conditions for fully apt judgment despite being negligent. Yet full-fledged HEV excludes negligence, one might think.

One might respond that sensitivity to objective risk is simply more important from

---

<sup>28</sup>See esp. Sosa (2015: 82).

<sup>29</sup>See Sosa (2015: 107-8).

<sup>30</sup>Sosa (2017: 78), parenthetical italics mine.

<sup>31</sup>Note furthermore that the mere fact that Sosa sometimes describes the guiding state as a kind of 'awareness' doesn't imply that he is building in something with the phenomenology that is intuitively missing in the clairvoyance case. Apart from the fact that this would not help to explain why such awareness is needed for full skillfulness (which is intuitively implausible), Sosa (2015: 79) is explicit in a footnote that the second-order 'awareness' that guides one's affirmation to aptness isn't meant necessarily to involve conscious phenomenology: '[T]his awareness need be neither conscious nor temporally prior.'

the epistemic point of view than sensitivity to subjective risk, and hence see little cost in ignoring this kind of negligence. Here one might take inspiration from Arpaly (2003) and Harman (2011), who both suggest that one can be morally praiseworthy despite acting contrary to one's beliefs about rightness. And there are epistemic analogues of Arpaly's celebrated Huck Finn case that suggest as much. Someone might be wrongly persuaded by skeptical arguments that she ought not to take any experiences at face value, but continue to respond to what are in fact perfectly good perceptual reasons nonetheless. But while such cases do suggest that one can be all-things-considered praiseworthy despite ignoring subjective risk, they do not suggest that there is nothing of value in responding to subjective risk. Indeed, Arpaly (2003: 36) agrees that there is some value here. A fully conciliatory picture would be able to explain this residual intuition, and the apparatus of performance epistemology doesn't explain this intuition as a special case of any more general truisms about skill. Manifesting commitment to value in acting conscientiously is at least a somewhat fitting attitude to value, though it might be yet more fitting to respect value in the sense of 'respect' compatible with Huck Finn's respecting the value of Jim's humanity.

## 4.2 Enlightened Belief at the First Order

There is a further way in which UP and its relatives may prove unhelpful even if true.<sup>32</sup> UP promises to explain what we admire in subjects who have suitable second-order perspectives on their abilities. But it is not only such subjects whose beliefs exhibit HEV of the kind coveted by the internalist tradition.<sup>33</sup> There is a difference in HEV between subjects who believe for good epistemic reasons and subjects who don't. Believing for a good reason doesn't require ascending to the second order. There are unselfconscious subjects who believe for excellent reasons, as well as self-conscious subjects who don't deploy their self-consciousness whenever believing for good reasons. UP gives no explanation of the enlightened status they achieve.

Now, Sosa (2015) considers some of these worries. In response to the worry that his view cannot explain why believing for good reasons is better than believing for no reason, he makes two moves: (i) he denies the intuition that beliefs supported by reasons are better than beliefs not supported by reasons, and (ii) he suggests that he could accommodate this intuition because "judgment is always based on (rationally motivated

---

<sup>32</sup>Kelp (2014) makes a point kindred to the following, albeit about an earlier edition of Sosa's view.

<sup>33</sup>Admittedly, some major figures in the internalist tradition emphasize the higher order. But mentalists and acquaintance internalists do not (or need not); cf. Fumerton (1988, 1995).

by) some degree of confidence. . . .”<sup>34</sup> Besides saying that he doesn’t share the intuition that beliefs based on good reasons are better than beliefs based on none, he gives an example, writing: “I find it implausible that intuiting axioms. . . must somehow be inferior to believing theorems based on the premises provided by the axioms.”<sup>35</sup>

But these replies are not fully convincing. One needn’t needn’t take believing theorems on the basis of axioms to be the paradigm case of believing for reasons. One could hold that one also believes for reasons when one rationally believes the axioms. What reasons? The only options Sosa (2015) considers are mental states. But one might reject mentalist ontologies of reasons and prefer to view reasons as being, in the normal case, facts about reality. On one attractive view, for example, the faculty of Reason puts us into touch with abstract facts like the sheer fact that  $1 + 1 = 2$ . It is then on the basis of the “fact made manifest”, as McDowell would say, that one believes—or perhaps on the basis of *apparent* facts made *apparently* manifest, in the case of mere rationality. Facts made manifest are perfect reasons to believe that these very facts obtain.<sup>36</sup> Alternatively, one might say that Reason puts us into touch with the properties of abstracta, and it is on the basis of the properties made manifest that one thinks what one does.

Anyway, the problem needn’t turn on reasons. Part of the problem is that the only form of sensitivity<sup>37</sup> to risk that bears constitutively on full aptness in Sosa’s sense is sensitivity to objective risk. A person can exhibit such sensitivity while ignoring subjective evidence-relative risk, even radical such risk. Although in many realistic cases, decreased sensitivity to subjective evidence-relative risk interferes with affirming in the endeavor to affirm aptly, by decreasing one’s sensitivity to objective risk, this pattern does not hold of necessity. Even if one is doubtful about the importance of reasons in epistemology, one should not scoff at the intuition that non-negligent belief is better from an epistemic point of view than negligent belief. But it requires sensitivity to subjective evidence-relative, not just objective, risk.

---

<sup>34</sup>Sosa (2015: 108).

<sup>35</sup>Sosa (2015: 107).

<sup>36</sup>Believing that *p* on the basis of the evident fact that *p* needn’t be circular, as I note in Sylvan (2016: 366): ‘Note that believing that *p* on the basis of the fact that *p* seems fine in some cases. What, for example, justifies my belief that I’m in pain? Plausibly, the sheer fact that I’m in pain. Even if one does not adopt this model, it is a natural model.... This example suggests a deeper response. Not all responses to reasons are inferential. When one comes to believe one is in pain upon encountering that fact, one forms the belief in direct, non-inferential response to the fact. One doesn’t move in a circle because one’s process of belief-formation lacks inferential steps. If a fact can be one’s reason in this way rather than via inference, [we] needn’t permit moving in a circle.’

<sup>37</sup>I use “sensitivity” in an ordinary dispositional sense, not Nozick’s. Cf. Briggs and Nolan (2008).

## 5 Concluding Remarks

Let's take stock and reflect on the broader implications.

I defended a twofold conclusion. Firstly, I argued that if beliefs are performances or interestingly performance-like, we cannot use a principle like UP to explain HEV. Since such principles underpin the explanation the new performance epistemology offers, it hasn't succeeded in reducing HEV. Secondly, I argued that even if UP captured a truth about skillfulness or aptness, it would still not help to explain HEV.

One might wonder whether a conclusion as sweeping as (T) is justified by the reasoning in this paper. After all, (T) targets performance epistemology in general, not just Sosa's version of performance epistemology. Two alternatives to Sosa that come to mind would involve an appeal to (i) knowledge as the aim, or (ii) some explicitly reason-involving condition as the aim (e.g., compliance with reasons). But I don't think switching to these versions will promise more illumination unless they cease to be versions of performance epistemology.

Take the knowledge-first proposal first. Knowledge, one might think, is easy enough to be consistent with the lack of HEV. Like Foley (1987), I am happy to attribute knowledge to believers who are internally like Norman the clairvoyant and Truetemp, though like Lyons (2009), I suspect etiological conditions beyond aptness are needed. As long as this familiar kind of view is not excluded from the get-go, it won't be an automatic upshot of manifesting a skill to know that one's belief exhibits HEV. After all, mere animal knowledge can be a manifestation of a skill to attain animal knowledge. Nor will the core dialectic change if one deploys a new version of UP instead of squeezing HEV out of the aim of belief. Viewed as a deed or attempt, a performance that manifests a skill to achieve rather than a skill to succeed will not be improved by being guided by a belief that manifests a skill to know. Viewed as an aiming, the performance could be improved. But it couldn't constitute high-grade knowledge, and so some HEV will remain unexplained. In this way, the core dilemma remains.

What of shifting to a view on which belief aims at something that clearly entails compliance with reasons? Two versions of this proposal are worth separating. One takes compliance with reasons as primitive, while the other takes compliance with reasons to itself be a special case of the AAA pattern.

Take the second first. The obvious way to develop this view would be to hold that complying with a normative reason R in believing p consists in having a motivating reason that corresponds to R as a manifestation of one's skill to select operative reasons

that correspond to normative reasons. Sosa has himself experimented with this view in collaboration with me.<sup>38</sup> But such a view won't help much here. For understood in this performance-theoretic way, compliance with reasons might itself be insufficient for HEV. Consider a superbindsighter who blind-sees that p, and then forms the belief that q as a manifestation of a skill to form the beliefs that are supported by p-like reasons. Assume p entails q. According to the account at issue, this superbindsighter complies with the reason given by the fact that p, by conforming to that reason as a manifestation of reason-conforming skill. But her resultant belief doesn't attain HEV.

A different view would take compliance with good reasons as primitive.<sup>39</sup> Adopting this view here seems to concede defeat, by admitting that some non-performance-theoretic normative property is at least as fundamental as the performance-theoretic ones. And it concedes defeat too early. It is no less plausible that compliance is analyzable as conformity with reason through skill than it is that knowledge is analyzable as correctness through skill. Anyhow, it remains unclear that compliance with reasons yields HEV. Compare Huck Finn, who is sensitive to the reasons given by Jim's humanity, but in an importantly blind fashion. One could switch gears again and appeal to *reflective* compliance with reason. But now we invoke what we sought to analyze.

I conclude that it won't help to invoke richer aims. Appealing to knowledge doesn't help. Animal knowledge that manifests a skill to gain animal knowledge needn't entail HEV. Appealing to compliance with reasons doesn't help, since the kind of compliance with reasons that performance epistemology can illuminate doesn't entail HEV unless HEV is built into the underlying picture of reasons or of reasons-responsiveness.

Now, for all I have argued, there may be a more general kind of evaluation of which skill evaluations of acts and different evaluations of attitudes might each be special cases. The more generic term 'competence' is suggestive. Indeed, I suspect we should distinguish between competence in the ordinary broad sense and skill. But as I stressed at one point in the paper, that conclusion bespeaks no victory for performance epistemology. For this broader kind of competence is not the property of performance epistemology. Framing our questions with the help of this property won't anyway help illuminate HEV in the way Sosa (2015) suggests. To take skill evaluations of actions as a special case, this

---

<sup>38</sup>See Sylvan and Sosa (2017).

<sup>39</sup>One could also impose more demanding subjective constraints on compliance with good reasons. But to do so would likely involve abandoning pure performance epistemology. The account of compliance that is structurally analogous to the account of knowledge is one on which compliance consists in success—i.e., having an operative reason that corresponds to a normative reason—which manifests a skill to succeed so understood.



more general form of evaluation must be one for which UP-like principles don't automatically hold. There will remain substantial differences between the theory of performance normativity and the theory of attitude normativity even if each can be framed as special cases of something more general (e.g., a theory of competence).<sup>40</sup>

## References

- Arpaly, N. 2003. *Unprincipled Virtue*. Oxford: Oxford University Press.
- Baehr, J. 2012. *The Inquiring Mind*. Oxford: Oxford University Press.
- Block, N. 1995. 'On a Confusion about the Function of Consciousness.' *Brain and Behavioral Sciences* 18: 227–247.
- BonJour, L. 1980. 'Externalist Theories of Empirical Knowledge.' *Midwest Studies in Philosophy* 5: 53–73.
- Briggs, R. and Nolan, D. 2012. 'Mad, Bad, and Dangerous to Know.' *Analysis* 72: 314–316.
- Chrisman, M. 2012. 'The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions.' *Journal of Philosophy* 588–612.
- Foley, R. 1987. *The Theory of Epistemic Rationality*. Cambridge MA: Harvard University Press.
- Frankfurt, H. 1971. 'Freedom of the Will and the Concept of a Person.' *Journal of Philosophy* 68: 5–20.
- Fumerton, R. 1988. 'The Internalism/Externalism Controversy.' *Philosophical Perspectives* 2: 443–459.
- Fumerton, R. 1995. *Metaepistemology and Skepticism*. Lanham: Rowman & Littlefield.
- Geach, P. 1956. 'Good and Evil.' *Analysis* 17: 33–42.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Harman, E. 2011. 'Does Moral Ignorance Exculpate?' *Ratio* 24 443–468.
- Kelp, C. 2013. 'Knowledge: The Safe-Apt View.' *Australasian Journal of Philosophy* 91: 265–278.
- Kelp, C. 2014. 'Unreflective Epistemology.' *Episteme* 11: 411–422.
- Korsgaard, C. 2008. *The Constitution of Agency*. Oxford: Oxford University Press.
- Korsgaard, C. 2009. *Self-Constitution*. Oxford: Oxford University Press.
- Lyons, J. 2009. *Perception and Basic Beliefs*. Oxford: Oxford University Press.
- Miracchi, L. 2015a. 'Competence to Know.' *Philosophical Studies* 172: 29–56.
- Miracchi, L. 2015b. 'Knowledge Is All You Need.' *Philosophical Issues* 353–378.
- Peacocke, C. 1992. *A Study of Concepts*. Cambridge, MA: The MIT Press.
- Pritchard, D. 2012. 'Anti-Luck Virtue Epistemology.' *Journal of Philosophy* 109: 247–279.
- Ridge, M. 2013. 'Getting Lost on the Road to Larissa.' *Nous* 47(1): 181–201.
- Sosa, E. 2007. *Apt Belief and Reflective Knowledge: A Virtue Epistemology (Volume I)*. Oxford: Oxford University Press.
- Sosa, E. 2010. *Knowing Full Well*. Princeton: Princeton University Press.
- Sosa, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Sosa, E. 2017. *Epistemology*. Princeton: Princeton University Press.
- Sylvan, K. 2016. 'Epistemic Reasons I: Normativity.' *Philosophy Compass* 11(7): 364–376.

---

<sup>40</sup>I thank Bob Beddor, John Greco, Chris Kelp, Elseltijn Kingma, Matthew McGrath, Conor McHugh, Lisa Miracchi, Simon Rippon, Mona Simion, Ernest Sosa, Ralph Wedgwood, Daniel Whiting, Fiona Woolard, two anonymous reviewers for this journal, and others in audiences at KU Leuven and the 2017 Bled Epistemology Conference for helpful comments on this paper.

- Sylvan, K. and Sosa, E. 2017. 'The Place of Reasons in Epistemology' in Star, D. (ed.) *The Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Vargas, M. (ed.) 2016. *Performance Epistemology*. Oxford: Oxford University Press.
- Watson, G. 1975. 'Free Agency.' *Journal of Philosophy* 72: 205–20.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.