Synthese (2020) 197:5301–5317
https://doi.org/10.1007/s11229-017-1510-0

S.I. : THE EPISTEMOLOGY OF ERNEST SOSA

Functional belief and judgmental belief

Kate Nolfi[1]

Received: 23 January 2017 / Accepted: 24 July 2017 / Published online: 2 August 2017

Abstract A division between functional (animal) belief, on the one hand, and judgmental (reflective) belief, on the other, is central to Sosa's two-tier virtue epistemology. For Sosa, mere functional belief is constituted by a first-order affirmation (or, perhaps, a simple disposition to affirm). In contrast, a judgmental belief is an intentional affirmation; a performance which is partially constituted by the believer's endeavor to affirm truthfully, and reliably enough. If, *qua performance*, judgmental belief is like the hunter's shot or the baseball player's swing, mere functional belief is much more like a heartbeat. This paper explores whether we should accept Sosa's distinction between mere functional belief and judgmental belief, and, if we should, how recognizing this distinction ought to shape our epistemological theorizing. Accordingly, the first aim of this paper is expository. It is to further clarify Sosa's contrasting categories of functional belief and judgmental belief and to attempt to characterize explicitly the role that the division between functional belief and judgmental belief plays in Sosa's two-tier virtue epistemology. The second aim of this paper is more critical. It is to articulate and begin to evaluate a series of concerns regarding whether Sosa's division between functional belief and judgmental belief is well-founded, and so to explore whether a virtue-theoretic performance epistemology ought to embrace the sort of two-tiered account of cognitive performance that Sosa favors.

Keywords Belief · Cognitive performance · Judgment · Reflection · Virtue epistemology · Epistemic evaluation · Cognitive agency

✉ Kate Nolfi
knolfi@gmail.com

[1] The University of Vermont, 70 South Williams Street, Burlington, VT 05461, USA

1 Introduction

The foundational commitment of Sosa's distinctive approach to epistemological theorizing—Sosa's revolutionary insight—is that beliefs (and other cognitive attitudes) are performances, and so epistemic normativity is a species of performance normativity. A performance, for Sosa, is a doing with a constitutive aim. And this aim structures the set of norms or standards with respect to which we can evaluate the performance *as a performance of the sort that it is*. Accordingly, for Sosa, epistemic norms or standards *are just* those norms or standards that evaluate our cognitive performances as the kinds of performances that they are. And our cognitive performances are appropriately subject to epistemic evaluation by virtue of having the constitutive aim(s) that they have.

Sosa argues that performance domains exhibit a AAA normative structure: a performance is accurate or successful insofar as it achieves its constitutive aim. It is adroit insofar as it manifests the performer's competence to achieve the aim. And it is apt insofar as the performance's success manifests the performer's competence to perform in a way that achieves the performance's constitutive aim. When a performance is apt, and only then, the performer is appropriately credited for the performance's success.

If beliefs (and other cognitive attitudes) are performances, then the epistemic domain exhibits a version of this AAA normative structure. And for the sort of virtue-theoretic, performance-based approach that Sosa's work has brought to prominence, this result serves as the starting point for epistemological theorizing. For the virtue-theoretic performance-based epistemologist, spelling out the specific content and structure of epistemic norms, and so understanding the character of epistemic evaluation, requires spelling out the metaphysical nature of human cognitive performances as partially constituted by their constitutive aims.

There are, of course, a variety of different ways in which a virtue-theoretic performance-based epistemology might develop from here.[1] I restrict my attention in what follows to Sosa's own "two-tier" strategy. Sosa develops and defends a bifurcated account of human cognitive performances, according to which human cognitive performances come in two flavors. For Sosa, mere functional belief is constituted by a first-order affirmation (or, perhaps, a simple disposition to affirm). In contrast, a judgmental belief is an intentional affirmation; a performance which is partially constituted by the believer's endeavor to affirm truthfully, and reliably enough. If, *qua performance*, judgmental belief is like the hunter's shot or the baseball player's swing, mere functional belief is much more like a heartbeat. And since judgmental and functional belief are different sorts of cognitive performance, they are governed by correspondingly different constitutive norms, and subject to differing species of evaluation.

[1] The kind of virtue-theoretic, performance-based approach to epistemological theorizing that Sosa's work has pioneered and brought to prominence is, at least in its most general form, now fairly popular, and has been developed in a variety of different ways by a long list of different theorists (e.g. Greco 2010). For a useful introduction to the blossoming contemporary literature on this approach, see Fernández Vargas (2016).

This paper explores whether the proponent of a performance-based virtue epistemology should accept Sosa's distinction between mere functional belief and judgmental belief and adopt the corresponding two-tiered framework for understanding the character of epistemic evaluation. Accordingly, the first aim of this paper is expository. It is to further clarify Sosa's contrasting categories of functional belief and judgmental belief and to attempt to characterize explicitly the role that the division between functional belief and judgmental belief plays in the kind of two-tiered epistemological picture that Sosa favors. The second aim of this paper is more critical. It is to articulate and begin to evaluate a series of concerns regarding whether Sosa's division between functional belief and judgmental belief is well-founded. If the arguments of this paper are successful, then the virtue-theoretic performance epistemologist ought to reject the sort of two-tiered account of cognitive performance, and the resulting the two-tiered epistemological framework, that Sosa favors.

2 Distinguishing judgmental and functional belief

On Sosa's view, there are (at least) two modes through which a performance might have its constitutive aim(s). And, as a result, there are (at least) two senses in which, when all goes well, we appropriately credit a performer for a successful performance. First, a performance can have the constitutive aim(s) it has by virtue of the fact that executing the performance constitutively involves a kind of intentional aiming—what Sosa terms *endeavoring*—on the part of the performer. Thus, the bow hunter's performance in taking a shot is guided and/or partially constituted by her intention to kill her prey. And the basketball player's performance in taking a free throw is guided and/or partially constituted by her intention to put the ball through the hoop. These performances are fully agential: they involve a kind of intentional action. Paradigmatically, at least, performances of this sort are voluntary, under the performer's direct control, consciously undertaken, and freely chosen.

Alternatively, however, a performance's constitutive aim(s) may simply be grounded in and determined by the functional, teleological, or biological organization of the performer. A heartbeat—a performance of the circulatory system—aims to pump blood, but certainly not because the circulatory system intends to circulate blood by contracting the heart muscle. Similarly, the performances of a thermostat aim to modulate the ambient temperature, but not by virtue of the thermostat's intending to do so. Instead, both the thermostat and the circulatory system are constituted functionally (by human design, in the case of the thermostat, and biologically, in the case of the circulatory system) so as to have a constitutive aim. As such, these performances are non-intentional and involuntary.

Sosa's recent work employs this division to advance the thesis that human cognitive performances—i.e. beliefs—take two different forms: functional (animal) and judgmental (reflective). At least when we restrict our attention to paradigmatic exemplars of each, functional belief and judgmental belief constitute fundamentally distinct sorts of cognitive performances, at least from the perspective of epistemological theorizing. For Sosa, judgmental belief constitutively involves an intentional aiming (what Sosa calls an *endeavor*) on the part of the believer. So, judgmental belief is, like the

hunter's performance in shooting at her prey, fully agential. In contrast, Sosa proposes that functional belief's constitutive aim is grounded in and determined by the functional, teleological, or biological organization of distinctive psychological/cognitive mechanisms by which functional beliefs are regulated. Functional belief, then, is like a heartbeat or the thermostat's switching of the furnace on when the temperature of the room drops below 68 degrees. It is non-intentional, involuntary, and often subconscious.[2]

Consider the way Sosa describes judgmental belief in the following representative passage.

> Not just any disposition to affirm counts as [judgmental] belief… real [judgmental] belief that p is a disposition to affirm that p in the endeavor to affirm correctly on that matter, with truth, where the extent to which the subject affirms in that endeavor does suffice by itself to yield the affirmation. You really believe (judgmentally) all and only when you are disposed to affirm when your endeavor is epistemically pure and disinterested, and aims to attain truth reliably enough on the question at hand.[3]

This passage makes clear that judgmental belief constitutively involves a kind of intentional aiming, which Sosa terms *endeavoring*. Specifically, judgmental belief involves the believer's intentional aiming at (i.e. endeavoring to achieve) truthful affirmation, attained reliably enough. Elsewhere, Sosa clarifies that judgmental belief involves an endeavor to affirm not only truthfully, but, indeed, aptly.[4] Thus, judgmental belief is constituted (at least in part) by the believer's intentional aiming not simply at truthful affirmation, but additionally at apt affirmation, wherein the truth of the affirmation's content manifests the subject's competence to affirm truthfully. And so judgmental belief aims at apt affirmation—truthful affirmation that manifests the subject's competence to affirm truthfully—by virtue of the fact that the believer's intention to so affirm guides and/or partially constitutes her believing.

Judgmental belief need not be conscious, deliberate, or voluntary, although paradigmatic cases often exhibit these features. Still, Sosa suggests that judgmental belief, as a kind of cognitive performance, is fully agential by virtue of being guided by the believer's intentional aiming. For Sosa, then, judgmental belief is simply a kind of intentional action. And this fact helps us understand both how and why judgmental belief is (at least paradigmatically) motivated by and based on reasons.

[2] Metaphysically speaking, it may be most appropriate to characterize judgmental belief and functional belief as lying on a spectrum and, accordingly, the extent to which any given cognitive performance manifests the performer's agency might be a matter of degree. Indeed, Sosa sometimes introduces the categories of judgmental belief and mere functional belief in a manner that exhibits sympathy for this way of thinking—see especially the opening paragraphs of Chapter 9 in Sosa's (2015). If this way of thinking is right, then the boundaries of Sosa's categories will be vague. Crucially, the concerns about whether Sosa's distinction is well-founded that I raise in what follows are *not* concerns about whether Sosa's distinction can be clarified so as to ensure well-defined boundaries between the categories of judgmental and merely functional belief. Rather, these concerns are more precisely framed as concerns about whether it makes sense to carve up the spectrum of cognitive performances (and so to locate borderline cases) in the way that Sosa's categories of judgmental and functional belief do for the purposes of epistemological theorizing.

[3] Sosa (2014) p. 173. A slew of similar passages appear in Sosa's other work. See, especially, Sosa (2015).

[4] See, e.g., Sosa (2015), Chapter 3.

In contrast with judgmental belief, functional belief is non-intentional, implicit, and subconscious. Paradigmatically, at least, functional beliefs "guide everyday action entirely below the surface of consciousness… [they are] not accessed through a simple conscious response to the relevant 'whether' question… [and] it can take thought and analysis to pull them up."[5] Functional beliefs have their constitutive aim not by virtue of involving an endeavor (not, that is, by virtue of involving intentional aiming on the part of the believer), but rather by virtue of their particular functional role in the human mind. Functional belief, then, is "aimed at truth, or at representing accurately, and reliably enough… through psychological or biological teleology."[6] Put differently, functional belief has a constitutive aim by virtue of having a (biologically and/or psychologically determined) proper function.

In the first instance, then, Sosa distinguishes judgmental belief and functional belief by appeal to the different ways in which these two sorts of cognitive performances come to have a constitutive aim. The result is what we might term Sosa's bifurcated account of human cognitive performance. But the contours of Sosa's categories of judgmental and functional belief are complicated by the way in which Sosa's work has developed a link between judgmental belief and reflection. Sosa suggests that, while judgmental belief necessarily involves a kind of second-order reflective endorsement (even if this endorsement is implicit or tacit), merely functional belief is simply "inapt for proper reflective endorsement."[7] So, on Sosa's account, the distinction between judgmental belief and functional belief seems to track a distinction between cognitive performances that manifest bare first-order competence to affirm truthfully, and reliably enough, and cognitive performances that additionally manifest higher-order competences (e.g. a competence to asses the reliability of one's first-order competence in the situation at hand). It may be appropriate to differentiate versions of a Sosa-style account of cognitive performance here. Perhaps the category of judgmental belief, as intentional affirmation in the endeavor to affirm truthfully, indeed aptly, necessarily aligns with the category of cognitive performance which constitutively involves a kind of reflective (higher-order) endorsement. Much of Sosa's discussion in *Knowing Full Well*—in particular, Sosa's arguments for the thesis that judgmental belief is on a par with suspension as an intentional withholding—suggests sympathy for this version of the picture. However, Sosa sometimes seems to suggest an openness to the idea that we might make sense of a kind of higher-order endorsement already present (or at least possibly present) in cases of merely functional belief.[8] And to the extent that this is possible, the contrasting categories of judgmental and functional cognitive performance will overlap, but not in perfect isomorphism, with the contrasting categories of merely first-order and higher-order ("reflective") cognitive performance. Regardless, this much is clear: judgmental belief constitutively involves a kind of higher-order reflective endorsement. In contrast, functional belief may or may not involve the same, or some similar, kind of higher-order reflective endorsement.

5 Sosa (2015) p. 67.

6 Sosa (2015) p. 51.

7 Sosa (2014) p. 179. See also Sosa (2010).

8 See, for example, the discussion on p. 51 of Sosa (2015).

Recall that, on Sosa's way of thinking, epistemic norms or standards *are just* those norms or standards that apply to our cognitive performances by virtue of these performances being the kinds of performances that they are. And so if judgmental and functional beliefs are interestingly different sorts of cognitive performances, then epistemic evaluation of each sort of cognitive performance will reflect their differences. Sosa's work suggests two ways in which this is so. First, Sosa's discussion in *Knowing Full Well* suggests that judgmental belief can achieve a kind of epistemic distinction that is categorically unavailable to mere functional belief. Only judgmental belief can constitute what Sosa calls knowledge-full-well. And this is because judgmental belief constitutively involves a kind of higher-order reflective endorsement, while functional belief does not. Accordingly, functional belief need not fall short if it fails to manifest a second-order competence in assessing the reliability of a first-order competence. But a judgmental belief that fails to manifest higher-order competence is thereby epistemically faulty or criticizable in at least one respect.[9]

Additionally, and perhaps relatedly, Sosa suggests that epistemic evaluation credits a believer differently for her judgmental beliefs and for her functional beliefs. Epistemic evaluations of judgmental beliefs have the same kind of agential character as do evaluations of a hunter's shot. The believer is (at least potentially) responsible for the quality of her judgmental beliefs in precisely the same fully agential sense that the hunter is (at least potentially) responsible for the quality of her shot. In contrast, epistemic evaluations of functional beliefs do not credit the believer as responsible for her beliefs in the distinctive way that we credit an agent as responsible for her actions. In this respect, evaluation of functional belief is on a par with evaluation of a heartbeat or of a thermostat's performance in switching the furnace on in response to a drop in the ambient temperature. It is, at least in the first instance, not evaluation of an agent and her intentional act (understood as a manifestation of her agency), but rather evaluation of a (biological/psychological) mechanism with a certain proper function. Accordingly, judgmental belief renders a believer appropriately subject to reactive attitudes (e.g. blame, or the distinctive kind of praise—more than mere approbation—that contrasts with blame). But mere functional belief cannot. Just as it would be inappropriate to blame the thermostat for its malfunctioning, or to blame a patient for her arrhythmic heartbeat, it would be inappropriate to blame the believer for her cognitive system's incompetent performance when this performance is mere functional belief (i.e. non-intentional, subconscious, and, in some sense, outside the subject's direct control).

3 Canvasing sources of support for distinguishing judgmental and functional belief

Ought the performance-based virtue epistemologist accept Sosa's bifurcated account of human cognitive performance? It seems to me that there are two strategies a propo-

[9] Although this interpretation of Sosa's (perhaps evolved) view is complicated by some of the discussion in Chapter 3 of Sosa (2015), which suggests some willingness to concede that mere functional belief might achieve some kind of importantly analogous epistemic status to that of knowledge-full-well.

nent of a Sosa-style account of human cognitive performance might adopt to build a case in its favor. First, she might try to show that our best theories of the human mind and of the metaphysics of agency provide reason for thinking that Sosa's distinction between judgmental belief and functional belief is metaphysically well-founded. Alternatively (and perhaps additionally), she might try to show that Sosa's account of human cognitive performance garners support from the character of our epistemic practice. In particular, if our epistemic practice systematically treats functional and judgmental belief differently, then we would have at least *prima facie* reason to embrace a Sosa-style bifurcated account of human cognitive performances as normatively significant.[10]

Each of these strategies, however, faces challenges. Challenges to the first (metaphysical) strategy suggest that Sosa's way of drawing the distinction between judgmental belief and functional belief does not illuminate the metaphysical nature of its target. In essence, the worry here is that, even assuming Sosa is right in taking our cognitive states to be performances, these performances simply do not fit neatly into the categories of functional and judgmental belief.[11] The distinction between functional and judgmental belief fails to carve the relevant bit of reality at the joints. Challenges to the second (normative) strategy suggest that our evaluative practice simply does not track anything like a distinction between judgmental and functional belief, and for good reason. If our evaluative practice in the epistemic domain does not consistently subject functional beliefs and judgmental beliefs to systematically differential treatment, then, assuming that there is no reason to think our evaluative practice is error-ridden, there is no pressure to recognize and embrace a distinction between judgmental and functional belief in epistemic theorizing.

Of course, these normative and metaphysical challenges might well be construed as mutually reinforcing. If our evaluative practice in the epistemic domain serves as one source of defeasible evidence that can guide our theorizing about the metaphysics of the domain of human cognitive performances, then a normative challenge will inspire or reinforce a metaphysical challenge. And if our evaluative practice is grounded in and responsive to the relevant metaphysics, then we have every reason to suspect that a successful metaphysical challenge will generate and, indeed, underwrite, a successful normative challenge to Sosa's bifurcated account of cognitive performance and two-tiered framework for epistemic evaluation.

Alternatively, it might be that showing that either one of these challenges can be effectively neutralized or overcome would be sufficient to show that Sosa's bifurcated account of human cognitive performance is well-motivated and defensible. If the distinction between functional and judgmental belief is central to our evaluative practice, and if our evaluative practice, itself, has value, then perhaps we have reason to embrace the distinction regardless of whether it figures in our best theories of the metaphysics of human mental states. And if this distinction does figure in our best theories of mind,

[10] Indeed, it seems to me that Sosa himself finds motivation for drawing the distinction between functional and judgmental belief from both these sources. See, especially, Sosa (2015).

[11] The worry here is *not* simply that the line between judgmental belief and functional belief is vague or blurry. See note 2 for further clarification.

then insofar as our actual epistemic practice fails to respect the distinction, perhaps the lesson to draw is simply that our practice needs revision.

I won't try to adjudicate this debate about the relationship between metaphysical challenges and normative challenges here. And for this reason, I will develop a metaphysical challenge and a normative challenge to Sosa's particular two-tiered version of performance-based virtue epistemology as independent lines of resistance below.

4 Metaphysical challenges

A metaphysical challenge to Sosa's two-tiered view will charge that an appeal to the notion of endeavoring (i.e. endeavoring of the sort that, according to Sosa, is constitutive of judgmental belief and constitutively absent in cases of mere functional belief) will not draw the contours of the categories of judgmental belief and functional belief so as to render these categories mutually exclusive, as well as significantly and interestingly non-empty.[12]

In order to make the force of challenge palpable, first consider two of Sosa's own examples: the patient's reflexive kick when the doctor strikes her knee with a mallet, and the hunter's performance in taking aim and shooting at her prey. The reflexive kick is something the patient does, but it is not a manifestation of the patient's agency. The hunter's shot, however, is a paradigmatic manifestation of agency: it is freely chosen, voluntary, deliberate, consciously directed, and intentional. It is a clear case of endeavor: an intentional action, characterized by its distinctive, intentional aim. It is easy to see, however, a great deal seems to fall in the messy terrain that lies between the hunter's shot and the patient's reflexive kick.[13] Think of the mindless way in which the engrossed movie viewer consumes her popcorn, eating till the bag is empty, the long-time resident's habitual movement to flick on the light switch as she opens the front door to her apartment, the unreflective, unthinking way in which an absent-minded professor might navigate the familiar route from a classroom back to her office while absorbed in thought or conversation, or the subconsciously regulated automaticity with which the practiced operator of a manual-transmission vehicle shifts between gears. These are certainly all doings. Moreover, it seems clear that these doings manifest the doer's agency (although perhaps less cleanly or crisply than the hunter's shot).[14]

[12] I leave other ways of drawing the contours of these two categories to one side here.

[13] Sosa characterizes the primary occupants of this region as *functionings* in both his (2013) "Epistemic Agency" and again in chapter 9 of his (2015) *Judgment and Agency.* But the kinds of examples of agential action to which I appeal in what follows do not fit neatly into the three-part (mere doing; functioning; endeavor) framework to which Sosa appeals in these pieces. Instead, they appear to lie somewhere between what Sosa terms *functionings*, on the one hand, and endeavors, on the other. If this is right, then we have some reason to worry that Sosa's framework is similarly ill-equipped to account for many of our cognitive performances.

[14] Sosa is sympathetic—see specially Sosa (2013) and chapter 9 of Sosa (2015). Sosa does not, however, put forward an account that explains how/why certain doings that fall short of (paradigmatically conscious, deliberate, and voluntary) fully intentional endeavors manifest agency. One lesson I draw from the discussion that follows is that we have reason to think such an account (i.e. an account that explains manifestations of agency without appeal to the kind of intentional aiming that constitutes endeavoring) will be better-suited to illuminate the way in which our cognitive performances manifest agency.

Can the notion of endeavoring help to clarify why (or in virtue of what) these doings constitute manifestations of the doer's agency? Saying that the long-time resident's habitual flicking of the light switch, or the practiced driver's shifting from third to fourth gear, constitutively involves an endeavor can seem to over-intellectualize the kind of psychological activity that underwrites these doings. But perhaps the notion of endeavoring can be de-intellectualized so as to illuminate these intermediate manifestations of agency.[15] Regardless, a dilemma presents:

Horn 1: If the movie-goer's popcorn eating, the long-time resident's habitual flicking of the light switch, the practiced driver's shifting from third to fourth gear, and other similarly intermediate cases manifest agency without involving endeavoring, then we have no reason to think endeavoring will figure prominently in the overwhelming majority of our cognitive performances. It now seems that our judgmental beliefs, if we host any at all, must be very, very few and far between. And, accordingly, it seems we have some reason to think that the category of judgmental belief is of at most marginal importance from the perspective of epistemological theorizing.

Horn 2: If, however, the movie-goer's popcorn eating, etc. constitutively involve (now involuntary, subconscious, and unreflective) endeavoring, then so should we think do the vast majority of our cognitive performances. Accordingly, an initial worry is overcome: if the vast majority of human cognitive performances involve endeavors of the relevant sort, then judgmental belief is commonplace. But now one might reasonably worry whether a careful survey of our cognitive performances will expose the class of merely functional beliefs as vanishingly small, and so without real significance for epistemological theorizing.

Let us take stock. The charge here is, in essence, that Sosa's account of judgmental belief as intentional action over-intellectualizes our cognitive performances, while Sosa's account of functional belief as on a par with the heartbeat under-intellectualizes our cognitive performances. Instead, the objector suggests that, in the vast majority of cases, human belief is not a matter of full-blooded intentional action (on a par with the hunter's shot) and also not a matter of bare teleological/biological/mechanistic performance (on a par with the heartbeat or the patient's reflexive kick). I have already argued that an appeal to the notion of endeavoring won't neatly partition the wide range of action-performances into two classes. Moreover, even if there is an account of endeavoring in the offing that carves the terrain between the patient's reflexive kick and the hunter's shot in a way that yields defensible verdicts about the movie-goer's popcorn eating, etc., a further worry now presents as pressing. After all, metaphysics of mental agency might well fail to mirror the metaphysics of agency in action. Human cognitive performances might manifest agency in an importantly different, and perhaps more fundamental, way than the hunter's shot. And if this is right, then an account of endeavoring which helps to illuminate the manifestations of agency or lack thereof that we see in the movie-goer, the long-time resident, the practiced driver, and the absent-

[15] One way of construing the positive proposal in Dickie (2016) is as an attempt to provide a suitably de-intellectualized counterpart to Sosa's notion of endeavoring. The generally concessive tone of Sosa's (2016) response to Dickie's proposal suggests that Sosa may be open to this sort of alternative.

minded professor won't necessarily help to illuminate manifestations of agency in the domain of human cognitive performances.[16]

The crux of this metaphysical challenge, then, is that the overwhelming majority of our cognitive performances, manifest agency—i.e. they constitute our *doings*—in ways that go beyond the reflexive kick, but fall short of the hunter's shot. I am gazing out the window and believe that it is raining—perhaps I even absent-mindedly make some comment to this effect. I hear the phone ring, I look at the caller ID, and I believe that when I answer I will hear my partner's voice. As we get into the car to head for the theatre, I notice that it is already 7:40 and I remember that the movie starts at 7:45, so I believe that we'll miss the previews, the opening credits, and perhaps the beginning of the film. I happen to glance at the conference program and find myself believing that the last talk of the day ends at 5:30 pm. As a result I believe that I'll be able to keep my promise to meet a colleague at the hotel bar for a drink at 5:45 pm even if I stay through the end of the last session. Clearly these beliefs are not the product of explicit, conscious deliberation or reflection on the import of the available evidence. In this respect, at least, they are quite unlike the belief I might form as a member of a jury tasked with evaluating whether it is beyond a reasonable doubt that the accused is guilty. These everyday, commonplace, and mundane (hereafter abbreviated as "ECM") beliefs are not precipitated by my explicitly intending or endeavoring to affirm truthfully or aptly. ECM beliefs are formed without effort and without conscious direction. Indeed, they are beliefs that it seems one simply finds oneself holding. Still, these sorts of ECM cognitive performances are not comfortably characterized as merely reflexive, simply beyond the scope of our agency. Even if, given my circumstances, my ECM beliefs are ones that it seems I cannot help but have, they are, nevertheless, robustly *mine*. *I* deserve a kind of credit in these cases for believing as I do. In contrast, *I* do not deserve the same sort of credit for the excellence of my reflexive response to the doctor's mallet (perhaps the credit goes to my neural circuitry). And perhaps even more tellingly, *I* am not appropriately credited for whatever epistemic virtues a belief I hold as a result of, e.g., hypnotic suggestion, happens to have. Moreover, it does not seem that I can disown my everyday, commonplace, and mundane beliefs in the way that I might disown the reflexive kick or the belief I hold as a result of hypnotic suggestion. For better or for worse, it seems that my ECM beliefs are beliefs for which I am answerable and appropriately held to account.[17] I am on the hook, so to speak, for these beliefs: they are the sorts of beliefs for which I am expected to have good reasons (even if, as it happens, I never end up articulating these reasons to myself or to others.) Indeed, this is especially vivid in cases where the believer expresses ECM beliefs (e.g. in giving testimony), or relies on these sorts of beliefs (e.g. as a basis for action or further reasoning). So, our evaluative practice in the epistemic domain

[16] Hieronymi (2006), among others, defends this sort of alternative view, according to which there is a distinctive and more fundamental sort of agency manifest in our mental attitudes. This sort of view can seem especially attractive if one focuses on the way in which the toxin puzzle and its analogues seem to expose belief, intention, and other mental attitudes as non-voluntary and impervious to direct control by the will.

[17] Of course, the mere fact that my belief is epistemically criticizable is not enough to ensure that I am epistemically blameworthy. Quite plausibly, I am only epistemically blameworthy if my belief is epistemically criticizable by virtue of manifesting some sort of incompetence.

is premised upon the assumption that these beliefs are *our* doings in a way that our heartbeats and our hard-wired reflexive responses are not.[18] And this is some evidence that our ECM beliefs somehow manifest our agency without being fully voluntary, deliberative, and/or reflective, like the hunter's carefully-considered shot.

All this suggests that a great many of our beliefs (perhaps along with a slew of other sorts of mental performances) may be genuinely agential without being fully intentional endeavorings aimed at apt affirmation. And so it seems that these performances manifest our agency in some other yet-to-be-specified way. If this is right, then Sosa's distinction between judgmental and functional belief is metaphysically inadequate to the phenomenon it aims to illuminate. These categories, at least as Sosa characterizes them, do not pick out two importantly contrasting and metaphysically distinct sorts of performances, both of which figure prominently in human cognition. And, accordingly, the kind of bifurcated account of the domain of human cognitive performance that divides our cognitive performances into two classes according to whether these performances involve endeavoring is misleadingly oversimplified and radically incomplete. It looks as though the sort of two-tiered framework for understanding the character of epistemic evaluation that Sosa grounds on this bifurcated account of cognitive performance is without a sound metaphysical foundation.

5 Normative challenges

Does Sosa's particular brand of two-tiered performance-based virtue epistemology fare better against normative challenges? Recall that a normative challenge aims to undermine Sosa's bifurcated account of cognitive performance and the epistemological framework grounded therein by appeal to the nature of our evaluative practice. The charge here is, in effect, that our everyday epistemic evaluations fail to respect anything like the distinction that Sosa suggests between judgmental and functional belief. If our evaluative practice in the epistemic domain routinely treats what Sosa would have us classify as functional beliefs as on a par with what Sosa would have us classify as judgmental beliefs, or if we routinely measure functional beliefs and judgmental beliefs against the same set of epistemic norms or standards, then, even if the distinction between judgmental and functional belief is metaphysically respectable, we have some (defeasible) evidence that the distinction is without real epistemological significance.

So, does our evaluative practice betray any kind of coherent and sustained commitment to treating judgmental and functional belief as normatively distinct? There is at least some reason to think that it does not.

To see why, consider an illustrative pair of cases. First, imagine Flora is driving home from work. The route is exceedingly familiar and Flora is deeply absorbed in thought (perhaps she is thinking about how to best structure the epistemology paper she is drafting) as she drives. At a certain point along the route, Flora changes lanes so that she'll be able to make a left turn at the next intersection. She automatically turns to check her blind spot, flips on the blinker, and then begins to guide the car from the

[18] I offer one account of that in virtue of which our beliefs manifest our agency (and so count as *our* doings) which is non-intentional, non-voluntary, non-deliberative, and non-reflective in Nolfi (2014).

center lane into the left lane. As a result of checking, Flora functionally believes that her blind spot is clear and so it is safe for her to change lanes. Moreover, this functional belief guides Flora's behavior: it is precisely because she so believes that she begins to execute the lane-change by turning the steering wheel. Still, it is easy to imagine that Flora's functional belief is wholly subconscious, unreflective, and unavailable for report. If a passenger were to have interrupted Flora's thoughts to ask whether the blind spot was clear just as Flora was about to turn her steering wheel, Flora would have felt uncertain and unable to answer sincerely either way. Such an inquiry would prompt Flora to re-check her blind spot, looking over her shoulder a second time, now in a consciously attentive and effortful way, before proceeding. But since Flora isn't faced with such an inquiry, she simply acts on her functional belief, and successfully changes lanes, all while deeply absorbed in thought about unrelated matters.

Now, compare Flora to Judy. Judy is a novice driver, still learning how to negotiate the city streets. Judy's driving instructor directs her to make a left turn at the upcoming intersection and reminds Judy that she will need to get into the left lane before making the turn. Judy consciously attends to the various steps of the task at hand as she turns to check her blind spot, flips on the blinker, and then begins to guide the car from the center lane into the left lane. As a result of checking, Judy judgmentally believes that her blind spot is clear and so it is safe for her to change lanes. Moreover, this judgmental belief guides Judy's behavior: it is precisely because she believes as she does that she begins to execute the lane-change by turning the steering wheel. Judy's belief is fully conscious and reflectively endorsed. Judy will easily and effortlessly answer in the affirmative when the instructor asks if she is *sure* that the path is clear for her to navigate into the left lane, and she could justify her confidence by citing the reliability of her visual experience if called upon to do so.

I stipulate that both Flora's belief and Judy's belief exhibit bare aptness. Still, Sosa's two-tiered account of the norms to which our cognitive performances are appropriately subject suggests that Judy's judgmental belief might achieve a special kind of high epistemic status: if all goes well, then Judy knows full well that her blind spot is clear. Flora's belief, however, isn't even a candidate for this sort of epistemic status. No matter how admirable, Flora's belief that her blind spot is clear won't amount to knowledge-full-well. Of course, since Flora's belief is a functional belief, its failure to achieve this kind of epistemic status is no shortcoming or defect. And so, given the kind of cognitive performance that each is, both Judy's and Flora's beliefs are beyond reproach, at least epistemically-speaking. Nevertheless, Sosa's account entails that Judy's belief will be epistemically praiseworthy in a way that Flora's is not (and, indeed, could never be). And so, crucially, it seems that Judy is admirable, at least *qua* epistemic agent, in a way Flora is not. Moreover, Sosa's bifurcated account of cognitive performance suggests that there is no dimension of epistemic evaluation along which Flora's belief is better (i.e. more admirable or more praiseworthy) than Judy's.

But these results do not jibe with the way in which our evaluative practice responds to the cases at hand.

First, notice that it can seem that Flora's belief is epistemically superior to Judy's, at least along one important dimension of epistemic evaluation. Once we bring the wider performative context into view, Flora's belief seems to manifest more compe-

tence, at least *qua* cognitive performance, than Judy's. Or, alternatively, the kind of performative competence that Flora manifests here seems to be more desirable (and so more admirable, more praiseworthy) than the kind of cognitive competence that Judy manifests. After all, Judy endeavors to develop, through continued practice, precisely the kind of cognitively effortless competence to determine whether a lane-change is safe that Flora already manifests. Flora's belief is expertly-formed in a way that Judy's is not. Yes, Judy reflectively endorses her belief. Flora does not and cannot. And perhaps this makes Judy's belief epistemically praiseworthy in a way that Flora's is not. Still, we'd feel safer with Flora at the wheel. And, crucially, it seems this is, at least in part, because we trust her to determine whether it is safe to merge (even while she is engrossed in conversation or wrapped up in her own thoughts) more than we trust Judy to do the same. Typically, at least, when Flora believes that her blind spot is clear and so that it is safe to merge, and she begins to turn the steering wheel in order to change lanes, we are content to simply presume that the beliefs on the basis of which she acts are competently formed—we won't ask Flora to affirm that her blind spot is clear because we take her epistemic competence for granted, and reasonably so. All this suggests that, although Flora's belief might not be categorically epistemically superior to Flora's, Flora's belief is epistemically praiseworthy in at least one way that Judy's is not. And it seems that, at least as it stands, Sosa's bifurcated account of our cognitive performances simply lacks the resources to vindicate this result.

There is, however, a line of response available to the proponent of Sosa's bifurcated account of our cognitive performances here that seems initially promising: the proponent of Sosa's bifurcated account might diagnose our temptation to regard Flora's belief as epistemically superior (along *any* dimension of epistemic evaluation) as a confusion. Of course, Flora's way of forming a belief about whether the path to merge left is clear might be better—i.e. more desirable and more admirable—from a practical perspective. Practically speaking, and perhaps even all-things-considered, Flora's cognitive performance is a *more appropriate kind* of cognitive performance, given her situation. Judy simply isn't yet capable of executing the relevant kind of cognitive performance—she is a novice, after all. And so she must execute a different, more cognitively demanding, kind of cognitive performance in order to accomplish the same practical end: she must form a judgmental belief that the path is clear to merge in order to safely maneuver her vehicle into the left lane. But epistemic norms and standards are not practical norms and standards. And so we should expect that, at least some of the time, epistemic and practical evaluation will diverge. What is practically admirable may be epistemically less so, and vice versa. The lesson here is meant to be that we must be careful not to interpret the ways in which Flora's functional belief is practically, or even all-things-considered, superior to Judy's as evidence that Flora's cognitive performance here is more *epistemically* praiseworthy or of higher *epistemic* quality. More generally, we might say that, speaking practically, there is a time and a place for mere functional belief. Indeed, there might be a host of different situations for which it makes sense and/or is desirable, from a practical perspective, to cultivate precisely the kind of unreflective automaticity that characterizes Flora's belief-formation. The reflective endorsement characteristic of Judy's judgmental belief can, at least in certain situations, seem unnecessary, not worth the extra cognitive effort, and therefore out of place. Nevertheless, when we take up the epistemic perspective, our focus shifts.

Epistemic evaluation is evaluation of our cognitive performances *qua the kinds of cognitive performances that they are*. Epistemic norms are just the distinctive constitutive norms of our cognitive performances. And as a consequence, the question of whether a subject engages in the most practically appropriate kind of cognitive performance, given her particular situation, is simply and straightforwardly external to epistemic evaluation.

But matters are not so straightforward. And so this line of response fails: we have good reason to reject the proposed diagnosis. To see why, consider a case from another performative domain: the marathon runner. In order to complete the race course in the shortest possible time, the marathon runner must carefully manage her energy reserves by moderating her pace throughout the race. Although her goal is to complete the race as fast as she can, she ought not run as fast as she can at each moment of the race. If she does, then she'll end up finishing more slowly (if she finishes at all) than she would have otherwise. Thus, runner A's performance on mile four might be superior along one important dimension of evaluation, even if slower, than another runner B's performance in running mile four. It is plausible that the explanation of this fact appeals to the way in which each runner's performance running the fourth mile of the race is constitutively embedded in or connected to the runner's other performances (on the fifth mile, on the ninth mile, over the course of the race as a whole, etc.). Because the runner's performance on the fourth mile is partially constituted—i.e. as the kind of performance that it is—by its relationships to these other performances, it is plausible that (at least some of) the constitutive norms that govern the runners' performances on the fourth mile will be sensitive to the level of competency that these performances manifest in calibrating the expenditure of limited physiological resources appropriately, given the demands of the constitutively connected performances that precede and follow the performance at hand. Similarly, it makes sense to think of our cognitive performances as constitutively connected. And, accordingly, it seems that, at least along one important dimension of epistemic evaluation, Flora's cognitive performance is especially epistemically praiseworthy—i.e. praiseworthy *qua the kind of performance that it is*—by virtue of involving an admirably calibrated expenditure of limited cognitive resources on the cognitive performance at hand. And more generally, it is plausible there are a great many similar situations in which a belief that seems to be automatic, unconscious, non-deliberative, unreflective, unintentional, effortless etc. it is epistemically better (i.e. more epistemically admirable or praiseworthy), at least along one important dimension of epistemic evaluation, than a belief that manifests the believer's explicit, conscious, and intentional endeavoring. And this result is one that Sosa's bifurcated account of cognitive performance and his resulting two-tiered framework for epistemic evaluation is poorly positioned to capture.

Moreover, there is a further feature of the way in which our evaluative practice handles Flora and Judy that stands in tension with Sosa's two-tiered framework. The way in which our evaluative practice credits Flora for her cognitive performance is on a par with the way in which it credits Judy for hers. In both cases, we are equally disposed to treat the subject as a genuine cognitive agent, fully responsible for and appropriately subject to reactive attitudes by virtue of her cognitive performance. And this suggests that our evaluative practice is premised upon the assumption that Flora's and Judy's performances are both genuinely agential. This fact is made vivid if we

imagine that either Flora's or Judy's belief that the blind spot is clear (and so that it is safe to merge left) is true, but incompetent. Perhaps the subject's visual ability to detect an object in the blind spot is somehow compromised, and that she might have been, but is not, appropriately sensitive to this fact, and so might have, but does not, compensate accordingly.[19] Regardless of whether we imagine Flora or Judy in the relevant set of circumstances, our inclination is the same: if we judge that it would have been easy enough for the subject to (either consciously, with attention and effort, or subconsciously, perhaps simply by turning her head a bit further) recognize and/or correct for her lack of competence, then we *blame* the subject for the fact that her performance was incompetent. Crucially, it doesn't seem that Flora gets off the hook, so to speak, for believing incompetently simply because her belief is merely functional and not judgmental—formed subconsciously, automatically, and non-intentionally. Thus, it seems that we regard Flora as responsible for her belief to the same degree and in more or less the same way that we regard Judy as responsible for her belief. Flora's belief is *her doing* as much as Judy's is *her doing.*

Perhaps it is worth considering another illustrative case. Imagine I believe that we have plenty of all-purpose flour in the pantry. Had I believed that we were out of all-purpose flour, I would have added all-purpose flour to the shopping list. But since I believe that we have plenty, I do not add all-purpose flour to the shopping list. My omission here need not be conscious, deliberate, or reflective—all-purpose flour just doesn't come to mind as a potential addition to the list (although if I had believed that we were low, it would have). If we are, in fact, out of all-purpose flour, and my belief is incompetent—perhaps as a result of my having mistaken a bag of whole wheat flour for all-purpose—then it seems I am rightly blamed for believing that we have plenty. And that I am blameworthy for my belief is part of what explains why I am blameworthy for leaving all-purpose flour off the shopping list. Crucially, it makes no difference if my belief is functional or judgmental. I am equally blameworthy regardless of whether (i) I find myself non-deliberatively, unreflectively, and quite unintentionally believing that we have a nearly-full bag of all-purpose flour after happening to glance at the relevant pantry shelf and seeing the bag of whole wheat or (ii) I consciously and deliberately, albeit hastily, contemplate the relevant pantry shelf in the endeavor to affirm aptly on the matter of whether we are stocked with plenty of all-purpose flour.[20] What matters is simply that I am fully capable of distinguishing a bag of whole wheat from a bag of all-purpose, that I happen to have failed on this particular occasion, and my failure was, in some sense, obviously avoidable. Had I gazed at the relevant bag for just a moment longer, I would have correctly categorized the bag as a bag of whole wheat

[19] We might appeal to the salient-to-our-subject possibility that there might be an appropriately located car draped in Harry Potter's invisibility cloak. Or we might appeal to loss of visual acuity at the periphery of the visual field, of which the subject is well aware, perhaps brought on some time ago by a very small stroke.

[20] My capacity to distinguish a bag of whole wheat from a bag of all-purpose does not generally require intentional effort, deliberation, reflection, or even conscious attention. And it is not as if I am *better* at telling bags of whole wheat from bags of all-purpose when adopt the goal of identifying the kind of flour bag before me as my intentional aim. Normally, I just *see* the bag as a bag of whole wheat or as a bag of all-purpose.

and, accordingly, I would not have formed the belief that we have plenty of all-purpose flour.

In general, then, the readiness with which we deploy reactive attitudes in the context of epistemic evaluation does not seem to track whether the belief that serves as our evaluative target is judgmental or functional (i.e. whether the belief constitutively involves an endeavor or not). Epistemic evaluation routinely treats functional beliefs and judgmental beliefs as rendering the believer appropriately subject to reactive attitudes in equal measure, and so seems to presuppose that functional and judgmental beliefs are equally agential performances. But this result is out of step with Sosa's bifurcated account of our cognitive performances and the resulting two-tiered framework for understanding the character of epistemic evaluation, according to which judgmental belief alone is genuinely, fully, or robustly agential by virtue of constitutively involving the believer's endeavor.

Of course, the proponent of Sosa's account might attempt to neutralize this charge by proposing that it is simply a mistake to classify—as I have in the preceding discussion—certain subconscious, non-deliberative, effortless, and apparently unintentional beliefs as a merely functional. On this line of response, although Flora's belief is wholly subconscious and unavailable for report, it nevertheless involves an endeavor (i.e. an intention) to affirm aptly. In fact, the proponent of Sosa's account might even cite those features of our evaluative practice highlighted above as some evidence that Flora's belief is an instance of judgmental belief. And she might bolster her case by pointing to the ease with which Flora might have been prompted to adopt a more reflective, self-conscious approach in an endeavor to affirm aptly regarding whether the blind spot is clear. For example, an inquiring passenger might have caused Flora to consciously take up and endeavor to answer the question of whether the path into the left lane is clear. When the passenger asks whether Flora is sure that it is safe to merge, Flora can and will easily glance over her shoulder again and re-check her blind spot. And, as a result, she will affirm—this time consciously, deliberately, voluntarily, and in an endeavor to affirm aptly—that the path is clear (and before re-checking, she will now withhold judgment on the matter). Perhaps the mere possibility that Flora could—and, indeed, would, if the situation demanded as much—bring her reflective capacities to bear in an endeavor to affirm aptly regarding whether the blind spot is clear is sufficient to render Flora's actual (automatic, subconscious, and non-reflective) belief judgmental rather than merely functional. But now the proponent of Sosa's bifurcated account finds herself back in the grips of a version of the metaphysical challenge discussed above. If Flora's belief is properly classified as a judgmental belief, then it is far from clear that there is anything like a substantial class of human cognitive performances that genuinely constitute mere functional belief.[21]

Finally, the proponent of Sosa's bifurcated account might, of course, simply accept the characterization of our evaluative practice that I have sketched above, and embrace

[21] Sosa suggests that "…what we learn about intentional belief—even conscious, intentional belief—should carry over to belief generally, whether intentional or merely functional" (Sosa 2015, p. 51). One way of putting my worry here is that, to the extent that the carry-over is seamless, isomorphic, and complete, the category of functional belief, as a distinctive kind of human cognitive performance, seems to dissolve.

the claim that our evaluative practice is, at least in certain spheres, simply mistake-ridden or error-prone. In effect, this rebuttal proposes that our everyday epistemic practice is in need of significant and substantive revision, and suggests that Sosa's bifurcated account of cognitive performances illuminates the sorts of revisions that are called for. But if we ought to favor less revisionary accounts over more revisionary competitors (at least *ceteris paribus)*, then the proponent of this line of response to the normative challenges I've raised here must show that no alternative account of human cognitive performance (and, in particular, an account which rejects a deep distinction between functional and judgmental belief) can mirror the various successes that Sosa's account enjoys. And this looks to be a tall order.

6 Conclusion

I have tried to advance two different sorts of challenges to Sosa's particular "two-tiered" brand of performance-based virtue epistemology. If the arguments above are successful, then we have good reason to think that Sosa's bifurcated account of our cognitive performances is not metaphysically well-founded and poorly positioned to vindicate certain aspects of our evaluative practice in the epistemic domain. More promising efforts to develop a performance-based virtue epistemology will involve giving up the sort of bifurcated account of our cognitive performances that Sosa favors and abandoning Sosa's corresponding two–tiered framework for understanding the character of epistemic evaluation.[22]

References

Dickie, I. (2016). Everybody needs to know? *Philosophical Studies*. doi:10.1007/s11098-016-0731-2.
Fernández Vargas, M. A. (Ed.) (2016). *Performance Epistemology: Foundations and Applications*. Oxford University Press: Oxford.
Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, *87*(1), 45–74.
Nolfi, K. (2014). Why is epistemic evaluation prescriptive? *Inquiry*, *57*(1), 97–121.
Sosa, E. (2010). *Knowing full well*. Princeton: Princeton University Press.
Sosa, E. (2013). Epistemic agency. *Journal of Philosophy*, *110*(11), 585–605.
Sosa, E. (2014). Epistemic agency and judgment. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms: New essays on action, belief, and assertion*. Oxford: Oxford University Press.
Sosa, E. (2015). *Judgment and agency*. Oxford: Oxford University Press.
Sosa, E. (2016). Replies to comments on judgment and agency. *Philosophical Studies*. doi:10.1007/s11098-016-0733-0.

[22] I trace the development of many of the ideas and arguments I offer here through conversations with Ernest Sosa, to whom I owe an immeasurable intellectual debt. I am also particularly grateful to Matt Weiner and two anonymous referees for *Synthese* for their comments on an earlier version of this piece.