Philos Stud (2013) 166:599–608
DOI 10.1007/s11098-013-0201-z

Risk sensitive animal knowledge

David Henderson · Terry Horgan

Published online: 6 September 2013

There is much that we find congenial in Ernest Sosa's *Knowing Full Well* (2011). We agree that knowledge is accurate belief that the agent attains by virtue of epistemically competent belief-forming and belief-maintaining processes. As Sosa puts it, for an agent to know is for the agent to attain accurate belief because the agent used adroit cognitive processes. This requires that, under the conditions that obtained in the episode, the agent's competent or adroit processes are causally explanatory of the agent's attaining true belief.

Central to Sosa's discussion is his distinction between animal knowledge and reflective knowledge. Sosa (2009, p. 135) characterized *animal knowledge* as "a variety of knowledge [apt belief] that does not require a knower to have an epistemic perspective on his belief, a perspective from which he endorses the source of that belief, from which he can see the source as reliable." Reflective knowledge, by contrast, does require such a perspective.

There are various ways one might construe what it is to have the sort of epistemic perspective on one's belief that qualifies the belief as a case of reflective knowledge. (Thus, the guiding notion of such a perspective is somewhat generic.) In *Knowing Full Well*, Sosa gives this notion the following explication, by way of explaining how he construes reflective knowledge: "Reflective knowledge … requires [in addition to what is required of a belief in order for it to qualify as an instance of animal knowledge] that the subject also believe aptly that his first-order belief is apt" (Sosa 2011, p. 11).

In what follows we will set forth some reasons to think that it would be preferable to explicate the pertinent notion of an "epistemic perspective" on one's belief in a

D. Henderson
Department of Philosophy, University of Nebraska-Lincoln, Lincoln, NE 68588, USA
e-mail: dhenderson2@unl.edu

T. Horgan (✉)
Department of Philosophy, University of Arizona, Tucson, AZ 85721, USA
e-mail: thorgan@email.arizona.edu

different, less intellectualist, manner than Sosa does—reasons prompted largely by what he himself says by way of motivating the putative need for a distinction between animal knowledge and reflective knowledge. We will propose an alternative explication, and we will explore some of its implications. One apparent implication is that the pertinent kind of epistemic perspective, when properly explicated, is already present in what Sosa calls animal knowledge. As a consequence, the distinction between animal knowledge and reflective knowledge apparently collapses: there is no adequate theoretical reason to posit some distinct form of knowledge over and above animal knowledge.

1 Motivating the distinction between animal knowledge and reflective knowledge

Sosa motivates his distinction via a guiding analogy: the hunter archer. (The book's cover illustration is a photo of an Early Classical Greek vase depicting Diana the huntress.) He writes:

> [T]he hunter archer's shot…can be apt in that its success, its accuracy, manifests the agent's competence in relevantly appropriate conditions (no wind, enough light, distance within proper bounds, and so on). But it, and its aptness, can also manifest the agent's meta-competence for target and shot selection. If so, then it is no accident that the shot is made in specific conditions where the archer's competence is up to the task of producing success with a high enough percentage. In other words, the agent's risk perception is then competent enough, and this competence is manifest in his knowledge that the level of risk is appropriate…. A performance is *fully* apt only if its first-order aptness derives sufficiently from the agent's assessment, albeit implicit, of his chances of success (and, correlatively), of the risk of failure. (pp. 10–11)

Shortly after setting out this characterization of a fully apt performance, as illustrated by the example of the hunter archer, Sosa draws upon it both to motivate and to characterize the distinction between animal and reflective knowledge:

> Animal knowledge is first-order apt belief. Reflective knowledge is animal knowledge aptly endorsed by the subject…. Knowing full well requires…that the correctness of one's first-order belief manifest not only the animal, first-order competences that reliably enough yield the correctness of the beliefs produced…. [One's] meta-competence governs whether or not one should form a belief at all on the question at issue, or should rather withhold belief altogether. It is only if this meta-competence is operative in one's forming a belief at all on that subject matter that one can reach the epistemic heights. (pp. 11–12)

Notice that here Sosa is fairly non-specific about the nature of the "meta-competence" that he says is required for reflective knowledge (and that he says is not required for animal knowledge). His own explication, as noted above, is in terms of the presence of a specific kind of *higher-order belief*: an apt belief that one's first-order belief is itself apt. But other potential explications could be given.

We ourselves find very plausible the idea that competent risk assessment, as an aspect of the process of forming a belief, is required in order for that belief to constitute fully human knowledge. But we doubt whether such competence needs to take the form of a *higher-order belief*; and we also doubt whether a first-order belief can qualify as any kind of knowledge if is formed in a way that *utterly lacks* the aspect of competent risk assessment. Let us elaborate.

2 An underlying functional distinction

As we have noted, much of Sosa's discussion of reflective knowledge supposes that the operative epistemic perspective must be managed at the level of (second-order) beliefs. However, at places, his discussion seems to suggest a way of thinking about what is important here—a functional role which, although it perhaps could be served by an epistemic perspective involving apt higher-order belief about first-order aptness, also could be served by other sorts of cognitive processes. His discussion often suggests that what is at issue might be functionally characterized as *cognitive sensitivity to epistemic risks*. We recommend that the epistemologically pivotal distinction be drawn in these terms. We urge that one distinguish *two functionally distinct capacities* in humans (and we think in many animals too)—and that one do so in a way that leaves empirically open how exactly the second function is managed as an aspect of belief formation. Ultimately we suspect that this distinction better captures what is needed for human knowledge. (It is possible that this is really Sosa's fundamental distinction, shorn of some incidentals with which it seems associated in his various discussions.)

On the one hand, there is something along the lines of what Sosa calls (first-order) *adroitness* in the generation of beliefs. This is an epistemic capacity for the generation of accurate beliefs. (The generated beliefs might be, but commonly are not, about risks.) As an epistemic *capacity*, this process-type possesses pertinent forms of reliability.[1] Call this **epistemic competence in belief generation**.[2] This is the functional idea that Sosa uses in characterizing his category of animal knowledge. Simple beliefs about things in one's near environment commonly serve as paradigmatic of beliefs generated via a process with this property. Yonder is a barn, a swallow, a dog, an intersection, … or simply, yonder is a red thing. However, the beliefs issuing from processes with this property also may be pretty theoretical and inferential: one may aptly form beliefs about plate tectonics as well as plates on the table, about bosons as well as bison. We agree with Sosa that in order to qualify as fully human knowledge, beliefs must be issue from processes that

[1] In keeping with Henderson and Horgan (2011), we would insist that the relevant forms of reliability include what we there call "transglobal reliability under suitable modulational control," and that this makes for processes that are globally (and commonly locally) reliable in epistemically hospitable global (and local) environments. We omit the details here.

[2] Perhaps this is a *class* of competences—such as the perceptual competence for generating true beliefs regarding North American wild flowers, the capacity for generating true beliefs about instanced computer malfunctions, and so on).

serve (intimately related) functional roles—roles that may not be apparent in the idea of epistemic competence in belief generation.

On the other hand, there is the capacity of being sensitive to epistemic risks. Call this **the capacity for epistemic risk sensitivity.** Sosa's analogy with capacities of the bow hunter is suggestive. The analogue to the epistemic competence for belief generation is the hunter's capacity to launch an arrow with a trajectory such that, barring unusual and unforeseen interference, it will likely hit the item at which one aims. The analogue to the agent's capacity for sensitivity to epistemic risks is the hunter's capacity for risk sensitivity in shot selection. The fully competent hunter, like the fully competent epistemic agent, is sensitive to information about conditions that might interfere (winds, obstructions, occlusions that might contribute to target misidentification, lighting, movement of the apparent target that might perhaps make for a more or less advantageous opportunity, chances for disturbance, and so on).

Sosa treats these two functional roles as separable in this strong sense: one's cognitive processes could serve the functional role of competent belief generation without serving in any way the functional role of epistemic risk sensitivity. This strong form of separability would follow, were one to accept his proposed explication of the capacity for risk sensitivity as being the capacity for forming apt second-order beliefs about the aptness of one's first order belief. Because he regards the two roles as so separable, and because of the greater cognitive sophistication that is presumably required in order to possess the capacity to form such second-order beliefs, he maintains that it is possible for a creature to possess animal knowledge that p without also possessing reflective knowledge that p.

We beg to differ—even though we do not dispute Sosa's over-arching conception of knowledge as accurate belief that is produced by epistemically competent belief-forming and belief-maintaining processes. Contrary to Sosa, we contend that a belief, in order to qualify as any sort of knowledge at all, must be produced and maintained by a form of competence that comprises *both* functional roles. Put another way, the kind of reliability that constitutes first-order adroitness must itself be a suitably risk-sensitive form of reliability.[3] We insist that a belief generated from cognitive processes of belief generation that are utterly devoid of sensitivity to extant indications of circumstantial risk of error is not any kind of knowledge.

We are not claiming, however, that in order to qualify as any sort of knowledge at all, a belief must exhibit all the features that Sosa himself calls "reflective

[3] The kind of reliability that we ourselves claim is constitutively required for epistemic justification, and hence also for knowledge—viz., what we call "transglobal reliability under suitable modulational control" (cf. note 2), is a form of reliability that is inherently risk-sensitive. Indeed, the apparent need to accommodate risk sensitivity within an adequate account of epistemic justification—and hence within an adequate account of knowledge too—was a principal motivating rationale for the form of reliabilism we advocate in Henderson and Horgan (2011). We call it *transglobal* reliabilism. Central here are processes of modulational control—such processes are responsive to information bearing on risks. In training up a normal human (or normal dog, monkey, or similar animal) as a perceiver of some class of phenomena, the human (or animal) subject's perceptual competence itself gets shaped so as to become automatically sensitive to a range of information associated with encountered risks. This kind of sensitivity via modulational control significantly enhances the reliability of the subject's perceptual processes and thus contributes to the subject's perceptual competence itself.

knowledge." In particular, we are not claiming that the belief that *p* must be accompanied by an apt meta-belief about the aptness of the belief that *p*. Rather, and again contrary to Sosa, we maintain that the pertinent notion of risk sensitivity—i.e., the pertinent explication of the idea of a belief-forming process that contains within itself a suitable "epistemic perspective" on the belief it spawns—is considerably less intellectualistic than Sosa's own proposed construal of risk sensitivity. Although suitable risk sensitivity must be inherent to the kind of first-order reliability that constitutes adroitness, it need not be—and in general, will not be—a matter of possessing an apt higher-order belief about the aptness of one's belief. (Indeed, even when such an apt higher-order belief is present alongside the first-order belief that *p*, its presence does not constitute possession of a new kind of *first-order* knowledge that *p*. Rather, given Sosa's over-arching conception of knowledge (which we ourselves find highly congenial), it constitutes the same *kind* of knowledge, but with a new *content*—viz., knowledge *that the belief that p is apt*.)

In Sosa's terminology, our contention could be formulated this way: genuine animal knowledge must be *risk-sensitive* animal knowledge, because the kind of reliability that constitutes first-order adroitness must itself be inherently risk-sensitive. (Hence the title of the present paper.) However, putting it this way perhaps suggests that Sosa is right to claim there is indeed a distinct kind of knowledge, over and above "genuine animal knowledge." We maintain otherwise: once the feature of suitable risk sensitivity is incorporated as a constitutive aspect of first-order adroitness—as it should be—there apparently remains no legitimate theoretical basis any longer to distinguish between two kinds of knowledge, animal and reflective.

In what follows we will set forth some considerations in favor of our claims in the preceding three paragraphs.

3 First-order risk sensitivity without epistemic meta-beliefs

Let us return to Sosa's guiding analogy: the hunter archer. It is very important to recognize that an adroit hunter might manage to exhibit functional sensitivity to risk in various ways. One way would be by articulately reflecting on (at least some of the) relevant risk factors, and then forming an explicit conscious belief about the likely adroitness and aptness of a potential shot (given those risk factors). Another way is too easily overlooked: the adroit hunter might have a trained capacity that manages to accommodate such information without articulation, automatically and quickly. The manifestation of this adroit risk sensitivity need not involve an occurrent *belief* about the risks—or any occurrent psychological state, either conscious or unconscious, that explicitly represents risk considerations. In cases that present no significant challenges, for instance, the archer's meta-level risk sensitivity need not induce any occurrent belief about the risk-relevant particulars to which it is sensitive. Perhaps the articulate representation of the scene by the hunter is in terms of the handsome large deer calmly browsing in the clearing, with the wind in the hunter's face—full stop. No occurrent representation of degree of risk here, and no need for one.

The cognitive case would seem parallel. Perhaps a biologist has a finely trained perceptual capacity to generate identifications of a range of birds: [male/female, adult/juvenile, of species x]. The capacity is clearly sensitive to a range of factors. Now suppose that the biologist also has a trained sensitivity to variations characteristic of a range of regional ecosystems—so that the biologist can also register common forest types (for example, Appalachian upland second-growth poplar-dominated forest, or old-growth Appalachian cove forest). This second kind of information is relevant to local expected base-rates for certain species. This information then might be accommodated in a capacity that is sensitive to the riskiness of attributing kinds with low regional base-rates. That is, the biologist might have in play processes that are sensitive to risks in identifying a critter in the local environment as a bird of a certain sort. The initial inclination to identify the critter moving in the yonder tree as a pine grosbeak might then be restrained by the agent's sensitivity to epistemic risks (a pine grosbeak would be weird in an Appalachian cove forest—unlikely to be correct). The initial inclination to identify the critter moving in another nearby tree as a summer tanager might occasion no such resistance from the agent's background sensitivity to epistemic risks.

A biologist with such an experientially fine-tuned capacity for forming bird-categorizing perceptual beliefs should surely count as *knowing full well*, when adroitly forming accurate perceptual beliefs about birds on the basis of this inherently highly risk-sensitive capacity—even when the biologist never forms any distinct, meta-level, beliefs about the in situ aptness of these first-order bird-classifying beliefs.

Turn now to a case involving the formation of a theoretical, non-perceptual belief. Suppose that someone engages in a bit of theoretical belief formation. Suppose that the agent is adroitly sensitive to many pieces of antecedent information, indeed all those pieces of information that constitute the data possessed. Consider an historical agent who judges that there are segments of the earth's crust floating on a dynamic molten mantle. The agent arrives at this belief first by noting that it would explain the range of geographical and archeological data on hand. The agent notes that alternative accounts then envisioned would not explain some ranges of the data. The agent then infers that this is how the earth's surface is composed. We may suppose that our historical agent here works at the level of data and what account would, if true, explain that data. Suppose that the agent's processes never lift their gaze to represent the form of inference deployed here, or the risks attendant to that kind of (abductive) inference, or the reasons why such risks are very largely absent in the current case. (Perhaps, if challenged, the agent would have trouble even characterizing such matters.) Still, the agent's processes would, or could be, pretty risk-sensitive *in their adroitness*, and the result would amount to a pretty significant form of theoretical inference. Intuitively, the agent could perfectly well know *full well*, without having formed any meta-level belief *at all* about the reasoning process whereby the belief was generated—let alone a meta-level belief to the effect that the first-order belief was aptly formed. Perhaps Sosa will insist that there is further, higher-order, knowledge to be had—and that methodologists and epistemologists could tutor our rather intuitive scientist into such knowledge. This we grant; and we grant too that such higher-order

knowledge could strengthen the scientist's first-order knowledge. What we deny, however, is that our scientist would thereby have acquired a new and different kind of first-order knowledge, viz., "reflective" first-order knowledge.

To get a further sense of how risk sensitivity can accrue automatically to suitably reliable first-order belief-forming processes, suppose that one has an agent who has a range of *experientially acquired* epistemic competences of the sort that would satisfy the requirement of epistemic competence in belief generation. We may suppose that these competences have been acquired by way of a training regimen in which the agent received feedback in a series of cases. For purposes of concreteness and breadth, we can imagine two cases:

> P: A case of perceptual competence in belief generation that has been acquired by training in cases with little articulate feedback. Perhaps the agent was merely scoffed at when having made an error, and given approving agreement when having formed a correct verdict. (Supposedly, in small scale horticultural societies much instruction, including that which is central to base-level concept acquisition, is largely the role of only marginally older children, and such teachers rarely give articulate feedback. In any case, suppose that that is how it went for our agent.) Suppose that our agent has acquired a pretty subtle capacity for perceptually gauging the matter in question—a capacity that is well attuned to the agent's actual local and regional geographical environment. Suppose it is a capacity to identify the instances of common plant-kinds. The cognitive processes that give rise to a belief that an item in the agent's environment is such-and-such a plant may automatically accommodate degrees and characters of occlusion, light levels, even of time of year, and much else.
>
> I: A case of reasoning competence of an inductive sort that has similarly been acquired by training in a group context in which certain specific inferences were scoffed at, and others were approved. When scoffed at, perhaps the agent's instructors mentioned a consideration that should have been given weight (or different weight), but they provided little systematic description of methods. The agents were there left to discern or acquire a sensitivity for themselves. In any case, our agent has acquired a pretty subtle capacity for gauging when an inductive inference would be risky. Perhaps, when judging an inference risky, the agent can sometimes articulate some respect in which it is so—but not always, and not in a way that is wonderfully revealing of just what systematically makes for various kinds of risks. Think of how graduate (or undergraduate) training manages to inculcate sensitivities that run beyond what can be (and doubtless beyond what is) systematically articulated in the training or by the agent.

Of either such agent, this much may be true in a particular case: by virtue of their trained-up process they each form beliefs and they each are sensitive to the epistemic risks in so forming those beliefs. However, in so doing, they *need not* form a belief about the adroitness of their processes, or the aptness of their belief. Commonly, fully competent human epistemic agents may roll along, forming perceptual belief after perceptual belief, or drawing inference after inference,

without generating any belief (apt or otherwise) about the aptness of the new first-order beliefs they issue. Nonetheless, in the very process of forming their perceptual belief, or drawing their conclusion, they may be sensitive to the risk. Indeed it is plausible that sensitivity to risk may be managed *in* the perceptual processes, or inferential processes, that they have acquired by long training.

4 No knowledge without risk sensitivity

We have been arguing that people can exhibit suitable risk sensitivity in belief formation—enough risk sensitivity for their first-order beliefs to count as instances of "knowing full well"—even if they do not form any meta-beliefs about the doxastic status of the pertinent first-order beliefs.[4] The pertinent risk sensitivity can be inherently present as an aspect of their first-order adroitness itself.

We also maintain that such first-order risk sensitivity, as an inherent aspect of the process whereby one forms and maintains one's first-order belief, is *required* in order for that belief to count as any sort of knowledge. There is no such thing as "pure animal knowledge"—i.e., knowledge without suitable risk sensitivity. In effect, this contention is quite common in the epistemological literature, even if it is not always formulated exactly this way. For, it is very common to contend that a true belief that was produced in a way that heavily depends upon *epistemic luck* fails (for that very reason) to qualify as knowledge; and the notion of an epistemically lucky belief is essentially just the notion of a highly risky belief that just happens to be true.

If we are right in claiming both (i) that the risk sensitivity required for knowing full well can be present without any meta-beliefs about the aptness of one's first-order belief, and (ii) that a belief cannot count as any sort of knowledge unless it is aptly generated by a *risk-sensitively* reliable first-order process, then the putative distinction between animal belief and reflective belief evidently collapses. There is just one kind of knowledge: viz. (given Sosa's over-arching conception of knowledge as apt belief) true belief that is produced by a risk-sensitively adroit first-order process, and is accurate because it is so produced. Of course, some of these are apt beliefs about apt belief.

It might be thought, however, that claim (ii) is false—false because there are beliefs that deserve to count as *some* form of knowledge (albeit not as instances of "knowing full well") even though they are not formed in a risk-sensitive way. If so, then there might still be a place for a version of Sosa's distinction between animal knowledge and reflective knowledge, even if one grants our contention (i). We will close by considering some remarks Sosa makes in the book which appear to lend support to the contention that claim (ii) is false. These remarks concern scenarios of the "fake barn country" variety. He writes:

[4] Or we have been suggesting this much. The real argument that there can be this kind of sensitivity in the processes of first-order belief generation is to be found in Henderson and Horgan (2000, 2011, Chaps. 7–8).

> Consider fake barn country. Better yet, switch for simplicity to a fake color environment. Someone views a red surface in good light and believes it to be red. What if all nearby surfaces that look red are actually white surfaces bathed in red light?…. Suppose our protagonist to be a color inspector brought in to determine the color of the surface of a plant, one of many throughout the world. Once a month he is whisked blindfolded into the high-security compound, and not only *is* not allowed but *would* never be allowed to enter the surrounding grounds (nor would surfaces from the grounds ever be allowed into the building)…. Does this fake color environment take away the subject's color competence? I cannot see that it does…. *Does he then know that the surface is red?*…. The color inspector knows he sees a red surface despite the many nearby fakes, provided he is protected against deceit [as in the case of the color inspector in the high-security compound]. Whether or not he knows the color of that surface depends on which surfaces nearby he might then easily have viewed. What accounts for this fact?…. [I]t is not plausible that competence, skill, or disposition is manifest at a certain location only if the host would have similarly succeeded elsewhere generally in the neighborhood…. We might therefore insist on a level of knowledge, animal knowledge, that is just apt belief, while suggesting that the subject in fake barn country might know on that level about a particular barn he then views…. Crucial to this approach is a distinction between such animal knowledge and a further level of knowledge, reflective knowledge. The latter, more demanding, requirement requires that the subject also believe that his first-order belief is apt, i.e., is one that manifests his competence. The inspector *might* satisfy this requirement, but only if the surfaces he might then easily have viewed would have been likely enough to be genuinely red. (pp. 82–92)

Nice try, we say, but not plausible. Moreover, there is a more plausible approach that invokes only a single kind of knowledge, the kind we have called "risk-sensitive animal knowledge." Let us take up these points in turn.

First, Sosa's approach to the case of the color inspector is implausible on its face. There is a perfectly appropriate sense, in context, under which the inspector (if he considers the matter) *does* aptly believe that his first-order belief manifests his competence. For, there is a perfectly appropriate sense in which it is just *vacuously* true that "other red-looking surfaces that he might easily have viewed would have been likely enough to be genuinely red"—since the folks running the compound take care to prevent him from being able to view those other red-appearing surfaces, thereby *preventing* them from being ones "he might easily have viewed." Indeed, this vacuously true construal of the modal locution 'might easily have viewed' seems to be the *most* appropriate construal, since the contextually salient reasons for attributing knowledge to the color inspector are so tightly tethered to the purpose for which he is there in the compound in the first place. More to the point, in the set of environments that the inspector can readily enter, he would be quite accurate in his color judgments; so there seems to be no good reason to think that the environments into which the inspector would never be allowed to enter would take away *either* his color competence *or* his competence to judge that his color beliefs are being aptly

formed. If aptness is understood in terms of a modal neighborhood, it seems that apt beliefs and apt beliefs about aptness may parallel each other.

Second, there is a more plausible approach to the case of the color inspector. It invokes just one notion of knowledge—viz., belief that is accurate, *risk-sensitively* (first-order) adroit, and apt—while also being modestly contextualist about how this univocal notion of knowledge should be applied in context. Two correlative ideas are central: (i) that suitable risk sensitivity is a constitutive aspect of first-order adroitness, and (ii) what *counts* as suitable risk sensitivity is a somewhat context-dependent matter. In ordinary (for philosophers!) scenarios of the "fake barn country" variety, the agent's belief counts as too dependent on luck—and thus as too *risky*—to qualify as knowledge. But one can concoct specific variations, such as Sosa's case of the color inspector in the high-security compound, for which what counts as a sufficiently non-risky mode of perceptual belief-formation does not depend, in any contextually pertinent way, on the various fakes that happen to be present nearby. Inherent risk sensitivity of the belief-generating process is still required for knowledge, in such a context, but the contextually relevant *kind* of risk sensitivity is independent of the presence of nearby fakes.

References

Henderson, D., & Horgan, T. (2000). Iceberg epistemology. *Philosophy and Phenomenological Research, 61*(3), 497–535.

Henderson, D., & Horgan, T. (2011). *The epistemological spectrum: At the interface of cognitive science and conceptual analysis*. Oxford: Oxford University Press.

Sosa, E. (2009). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 2). Oxford: Oxford University Press.

Sosa, E. (2011). *Knowing full well*. Princeton, NJ: Princeton University Press.