

Stephen R. Grimm, Fordham University

Draft: 6-30-14

Forthcoming in: *Performance Epistemology*. Ed. Miguel Angel Fernandez. New York: Oxford University Press.

THE VALUE OF REFLECTION

Human beings are distinctive in that unlike other animals we don't simply have beliefs and desires; we are also the kind of beings who can take a step back and evaluate or reflect on our beliefs and desires. We can do that, but why exactly would we want to? What is the good of it?

This, it turns out, is not so easy to explain. It is clear that the ancient Greeks held reflection in high esteem, for instance, but their reasons are unclear. Thus the Delphic Oracle tells us to "Know Thyself," but not why this reflective knowledge is a good thing for humans to have, or why a life without it would be diminished. And similarly Socrates tells us that "the unexamined life is not worth living," but not why an unexamined life is a lesser one, much less one not worth living.

A similar preference for reflection lives on in the epistemology of Ernest Sosa, and in his influential distinction between animal and reflective knowledge. For Sosa, animal knowledge comes about when one believes the truth as a result of one's intellectual abilities—for example, when you form a true belief that the cup is on the table as a result of your visual ability to pick out and identify cups. Beyond the animal, however, Sosa claims that human beings are also capable of knowing something in a reflective way. Thus we are capable not just of knowing that the cup is on the table (as an animal might), but of knowing this in light of our further knowledge that the source of our belief, in this case our vision, is trustworthy or reliable. Reflective knowledge is thus "enlightened knowledge" (Sosa 2009: 153). It is knowledge that does not just result in a mechanical way from our faculties but is instead guided by, and gains some sort of extra status from, our further knowledge that the source of our belief is reliable.

Although Sosa is an heir of Socrates in valuing reflection, his focus is also clearly different. For Socrates, the emphasis is on reflecting on one's own views regarding what is good or valuable or just. For Sosa, the emphasis is on reflecting on all of one's beliefs, and evaluating whether their sources are (or were) trustworthy. Further, while Socrates' judgment about the unexamined life is quite harsh—that it is not worth living—Sosa's view of "unexamined" or animal knowledge is considerably more generous. None of us, he says, would disdain acquiring the truth as a result of our animal abilities, just as none of us would disdain acquiring gold in the dark. But it is much better, Sosa claims, to acquire the truth while guided by a reflective perspective on our own abilities; this would

be like acquiring the truth in the light of day, where one can properly savor or appreciate it, or at least, where we can put to rest one's doubts.

Animal knowledge therefore is not epistemically bad, it is just that reflective knowledge is epistemically "better" (2007: 129; 2009: 193; 2011: 13); it deserves a certain "pride of place" (2004: 291) in our intellectual lives.¹ What I would like to do in this paper, however, is argue that the reasons that Sosa offers in defense of this view are unsatisfying. More exactly, I will argue that several of the reasons he provides do not in fact favor reflective knowledge over animal knowledge while others favor not just attaining a reflective perspective on one's animal knowledge, but towards a perspective on that perspective, and beyond. The question then becomes whether there is any defensible sense in which it is better to know reflectively as opposed to animally, and I will claim that there is such a sense, but it is not the one identified by Sosa: namely, that those who know reflectively exercise a kind of autonomy over their beliefs that those who know animally do not. Reflection therefore turns out to be valuable on this view not for what we would normally think of as "epistemic" reasons, but rather for something like moral or ethical reasons. In order to appreciate these points, however, we first need to get a better sense of what the distinction between animal and reflective knowledge amounts to in Sosa's work.

I. Animal vs. Reflective

According to Sosa, a belief amounts to animal knowledge just in case it is *apt*: that is, just in case the accuracy of the belief manifests the believer's competence at attaining the truth in this area (2011: 24). A bit of animal knowledge amounts to reflective knowledge, moreover, just in case that apt belief is aptly noted at the second level; that is, just in case the aptness of the object belief is guided by one's (apt) meta-belief that the source of one's object belief is reliable. To have reflective knowledge that my neighbor's dog is in the yard therefore requires not just that I aptly believe, as a result of my reliable eyesight, that the dog is in the yard, but also that this belief is guided by my apt second-order belief that my eyesight is reliable in these circumstances.

A few points of clarification. First, Sosa sometimes insists that one can have a second-order perspective on one's beliefs—and in particular, on the reliability of one's faculties—that is only implicit, or not brought to full awareness.² This complicates things, as Hilary Kornblith (2004) notes, because it is likely not just that humans but also other animals have some sort of implicit perspective on the reliability of their faculties—in which case my neighbor's dog might not just have animal knowledge that, say, a squirrel is in front of her, but

¹ Alternatively, the claim is sometimes that reflective knowledge is "higher" (2009: 151), or is "the best human knowledge" (2009: 153).

² E.g., "*Conscious* reflection on the spot is not required, however, since a second-order perspective can work beneath the surface of consciousness" (Sosa 2004, pp. 291-92).

reflective knowledge of this same proposition. Here I set this complication to one side: as I am understanding Sosa's thesis, in order to know reflectively one must actually reflect.³ This might seem obvious, but some of the subtleties of Sosa's view make it worth saying.

Second, there is one way in which the question of why reflective knowledge is better than mere animal knowledge has an easy answer—by my lights, *too* easy of an answer, and hence one which we should again set to one side.⁴ Thus one might claim that reflective knowledge that *p* is epistemically better than mere animal knowledge that *p* simply because there is more “epistemic stuff” on the scene when one reflectively knows—not just the first-order animal knowledge but the second-order reflective knowledge as well. Call this the “trivial answer” to our original question, one that claims that reflective knowledge is better simply because it constitutes more knowledge (cf. Sosa 2004: 285). I take it that this is not the right way to think about Sosa's claim because on his view it is the animal knowledge *itself* that is supposed to be improved by the addition of the perspective (see, e.g., Sosa 2011: 13). It is this claim, not the trivial one, which we are contesting.

II. The Main Argument

With these points in mind, Sosa's most common argument in favor of the betterness of reflective knowledge is that in the absence of reflective knowledge one's beliefs run the risk of incoherence.⁵ Suppose you are living contentedly with your animal beliefs but then begin to wonder whether your beliefs are justified or amount to knowledge. Unless you can answer in the affirmative, Sosa claims, your animal beliefs will then be seriously undermined. As he writes:

Nevertheless, a mere thermometer reaction to one's environment cannot constitute the best human knowledge, regardless of whether that reaction is causally mediated by experience. It is not enough that one respond to seeing white and round objects in good light with a 'belief' or 'proto-belief' that there is something white and round. Suppose one asks oneself 'Do I

³ As when Sosa talks of reflective knowledge where the source of one's belief is explicitly, reflectively “seen” or judged to be reliable (cf. Sosa 2004: 300). Two anonymous reviewers have pressed me on this point, arguing that the emphasis on belief over-intellectualizes Sosa's account, so let me add a bit more: If one supposes that the reflective monitoring can be highly implicit and “beneath the hood,” then the interest, and perhaps even the legitimacy, of Sosa's distinction between animal and reflective knowledge would be in danger. For it would then be hard to think of a case in which animal knowledge exists in the absence of reflective knowledge, because any animal that has this sort of (reliable) monitoring ability will count as having reflective knowledge—and that of course includes young children too. Whereas to begin with reflective knowledge seemed like a great accomplishment, because it allowed for properly human knowledge, rational defensibility, and the like, reflective knowledge would then almost come “for free,” once these monitoring abilities are on the scene.

⁴ Though towards the end of the paper I will again return to it.

⁵ See, for example, Sosa (2004: 302, 310; 2007: 115, 132; 2009: 146, 153).

know that this is right and round' or 'Am I justified in taking this to be white and round?' and one has to answer 'Definitely not' or even 'Who knows? Maybe I do know, maybe I don't; maybe I'm justified, maybe I'm not.' In that case one *automatically* falls short, one had attained only some lesser epistemic status, and not any 'real, or enlightened, or reflective' knowledge. The latter requires some awareness of the status of one's belief, some ability to answer that one does know or that one is epistemically justified, and some ability to defend this through the reliability of one's relevant faculties when used in the relevant circumstances. (Sosa 2009: 153)

In short, for those who consciously wonder about the status of their beliefs, Sosa claims that three options present themselves. (To narrow things down, suppose one is only wondering about some particular belief that *p*.)

Option A: One reflects on the status of one's belief that *p* and comes to think *p* is "definitely not" something one knows or is justified in believing.

Option B: One reflects on the status of one's belief that *p* and comes to think "maybe I do know that *p* and maybe I don't; maybe I'm justified, maybe I'm not."

Option C: One reflects on the status of one's belief that *p* and comes to think *p* is something one knows or is justified in believing.

Sosa unsurprisingly takes Option C to represent the person who has reflective knowledge. And his claim, again, is that Option C is definitely better than Option A and very plausibly better than Option B. Why? Elsewhere he makes it clear that,

Answer (a), and even answer (b), would reveal a certain lack of integration in that stretch of consciousness; only answer (c), of the three entirely avoids disharmony within that consciousness at that time. If one has to give answer (a), or even answer (b), one thereby falls short, and one's belief that *p* itself falls short. That belief is then not all it could be. (Sosa 2007: 115)

In other words, the idea is that Option C is better than Options A or B because only C displays intellectual coherence—only C "entirely avoids disharmony within that consciousness at that time."

What should we make of this argument? Notice, for one thing, that it is a much more limited version of the claim we began with, which seemed to hold, unrestrictedly, that reflective knowledge is better than animal knowledge. The argument here, by contrast, at best shows that *for those who reflect* (or wonder, or

express doubts about the status of their beliefs) reflective knowledge is better. The only way to return to the more unrestricted version of the claim, as far as I can see, would be not just to say “For those who reflect, reflective knowledge is better,” but rather that human beings *should* so reflect (or wonder or doubt). Or maybe that human beings are only properly satisfied, or can only bring inquiry to a proper close, when they have attained this sort of reflective endorsement. This would add an additional normative dimension to Sosa’s claim that it seems he needs, if he is to defend its unrestricted form.

In Sections 3 & 4 I will return to further thoughts Sosa offers along these lines, but it is worth pointing out now that even this more limited defense of the value of reflective knowledge does not obviously get him everything he wants. For even if we agree with what we might call the *meta-incoherence norm*—that for those who doubt, one has cognitive disharmony (or meta-incoherence) in the absence of reflective knowledge—it is nonetheless true that this norm can be satisfied in two different ways, both recognizably Sosian. Thus on the one hand you could reflectively endorse your first-order belief in an animally apt way: perhaps simply by affirming, at the meta-level, your Reidian trust in the reliability of your first-order faculties. On the other hand, you could reflectively endorse your first-order beliefs in a way that is backed by what Sosa calls a “structure of reasons” (2007: 121)—perhaps, as he suggests, reasons you acquire via cognitive science that support the reliability of your perceptual faculties.

It is clear however that Sosa thinks only the second sort of reflective knowledge, the one backed by the structure of reasons, is the one we really want. Why is that? Because, remember, the move to the meta-level in the first place was prompted by doubt that one’s first-order belief was reliably formed, and Sosa seems to think that only reasons can then dispel this doubt, not some sort of animal trust or faith.

Notice though how the argument has now shifted. What began with an appeal to the uncontroversial norm that meta-incoherence is to be avoided has now evolved in response to a more controversial norm—call it the *doubt norm*—along the following lines: that doubt can only be properly resolved by appeal to reasons. Once the doubt norm is on the scene, however, the main question is how it might be contained. If one accepts on the meta-level not just that sense perception is reliable, but that sense perception is reliable because of x, y, and z (where x, y, and z might be certain claims from cognitive science), then it is natural to wonder how claims x, y, and z might escape from doubt. But if they are subject to doubt, don’t we need reasons to support them as well? And so on.

In the following sections we will see that Sosa tries to evade the threat of regress by denying that doubt is cogent at the meta-level. Before we turn to this claim, however, it will help to ask why Sosa thinks there is normative pressure to doubt or reflect at the first-order level to begin with.

III. Other Norms of Belief

We have seen that for those who doubt the status of their beliefs, reflective endorsement of some sort seems necessary to stave off cognitive disharmony. But to return to our earlier question: What about those who fail to doubt or reflect? One can hardly say that their beliefs are meta-incoherent for the simple reason they have no meta-level beliefs that might fail to cohere. Or again, what about the person who fails to reflect not out of thoughtless neglect, but (as it were) willfully? The sort of person, that is, who declines to reflect because he thinks the activity is useless or idle or maybe even decadent.⁶ Suppose this person has animal knowledge that, say, his cup is on the desk, but when asked whether he ever wonders if his vision is accurate, simply replies “No thanks, I’m fine.” What if anything is wrong with this response?

One intriguing reply is that resting content at the animal level is inadequate because it somehow falls short of our natural goals as believers, and hence that there is something lacking or deficient about a belief that is not subjected to doubt or reflection. In this section I will consider two ways Sosa tries to support this claim.

The first is that as human beings, or perhaps as rational or intellectual beings, we have a natural desire to understand the world that leaves us unsatisfied with mere animal knowledge that *p*. Thus Sosa claims that when we consider what we care about from an epistemic point of view, “several epistemic values stand out” (2009: 136). To begin with, we want beliefs that are both (a) true and (b) safe. More than that, however, we would like:

(c) Understanding/explanation: often we would like not only to know a given thing, but also to understand it, to have an explanation. (And this leads to the next item.)

(d) Coherence: we would prefer that our minds not house a clutter of mere facts sitting there loose from one another. (2009: 136-37)

In another passage he likewise stresses the tight connection between understanding and coherence, along with their interlocking value:

[N]ote how the explanatory coherence of a body of beliefs comports with its providing the subject a corresponding yield in understanding. The cement that binds *B_p* and *B_q* in that body coherently would be a further belief of the form *B(p because q)* or *B(q because p)*. And such a belief

⁶ Here I leave out of consideration the still more complicated case where one might think, based on empirical work, that reflection is an unreliable process, one that is apt to lead us farther away from the truth. For more on this, see Tiberius (2010).

would embody some measure of understanding as to why p or as to why q. Therefore, the value of such coherence, explanatory coherence, is allied to the value of understanding. (2004: 312)

All told, the idea therefore seems to be that when we have mere animal knowledge that p, we have knowledge that satisfies our preference for the true and the safe, but not knowledge that satisfies all of our epistemic preferences. In particular, it fails to satisfy our preference to understand why p, a kind of understanding that brings with a correlative coherence among our beliefs (as opposed to a mere clutter).⁷

For the sake of argument let us grant there is something unsatisfying or incomplete about merely knowing that p, and that we naturally desire to understand why p as well (at least often, or characteristically, or insofar as we are human). What remains unclear, however, is how exactly this claim is supposed to support the value of reflective knowledge. Suppose I animally know that my coffee just spilled, and that as a human being I naturally desire to push beyond this knowledge in order to achieve an understanding of *why* it spilled (along with the correlative coherence this would bring). Surely what I would need in order to achieve an understanding of this sort is not some sort of perspective on the reliability of my perceptual faculties, but rather some knowledge of the “why” of p, or the cause of p: in this case, let us suppose, the knowledge that my coffee spilled because I bumped the desk with my knee. Notice however that this extra bit of knowledge is quite different than the sort of reflective knowledge that Sosa has in mind. For one thing, it is still an item of first-order knowledge, one that says nothing about the reliability of my faculties; for another, it is the sort of knowledge that is fully available to non-human animals and unreflective human beings. If the psychologists are right, animals at least from mice on up seem capable of grasping the causal dependencies that obtain in the world.⁸

In addition to the desire for understanding, a second reason why Sosa thinks there is normative pressure to move beyond the animal level is that as a rational human being I *should* doubt my first-order beliefs. Or at least, I should be concerned about the possibility of error among those beliefs. Thus Sosa approvingly quotes the ancient Pyrrhonian Sextus:

Let us imagine that some people are looking for gold in a dark room full of treasures.... [N]one of them will be persuaded that he has hit upon the gold even if he *has* in fact hit upon it. In the same way, the crowd of philosophers has come into the world, as into a vast house, in search of

⁷ As he reiterates elsewhere: “Understanding is intimately connected with coherence, since the explanatory interrelationship among our beliefs is bound to function as a, or even the, main component of epistemically efficacious coherence” (Sosa 2004: 298).

⁸ For more on this see Blaisdell (2006) and Taylor (2009).

truth. But it is reasonable that the man who grasps the truth should doubt whether he has been successful. (Quoted in Sosa 2007, p. 129; underline my emphasis)

What I *should* want then, is not just to believe that p in a way that I hope amounts to knowledge—I should want to believe that p in a rationally defensible way, or in a way that protects against doubt. And the move to the reflective level is apparently what it takes to resolve these doubts, or to make my belief defensible in the arena of reflection. Anything less would (or at any rate should) fail to satisfy me as a rational being.

What should we say about this? To begin with, we should resist the idea that doubt can only be removed, or rational defensibility can only be secured, by moving to the meta-level. Suppose that based on past experience I believe that my local bank will be open this Saturday. You then tell me that several local banks have been eliminating their Saturday hours in order to save money, leading me to doubt my belief. How might I resolve this doubt? Why not: by calling up my bank and asking whether they will be open on Saturday? Surely this strategy would be more natural and effective than taking some sort of meta-perspective on my faculties. But then it seems that doubts can legitimately be resolved without moving to the meta-level at all.

Or again, suppose my wife tells me that the car keys are near the toaster, that I believe her about this, but then I come to wonder whether the keys are really by the toaster, knowing that the kids sometimes move things around. How could I resolve my doubt? Here again, why not by going and taking a look? Then I will know for sure, and in a way that legitimately dispels my previous doubt—meta-level not required.⁹

Perhaps it will be said here that my doubts should not be so easily removed, or that rational defensibility is not so easily acquired. There is, after all, still the possibility that someone was playing a trick on me when I phoned the bank, or that I was hallucinating when I saw the keys near the toaster. But the more exotic the doubts become, the less obvious it likewise becomes that I “should” resolve these doubts. Arguably a rational human being could responsibly dismiss *these* doubts as idle, or exotic, or decadent, in a way that the previous doubts could not responsibly be dismissed.

There is also a more profound worry for Sosa, one that we gestured at in the previous section. Suppose the Pyrrhonians are right not just that we can

⁹ Perhaps what Sosa will say here is that I have now achieved a kind of (implicit?) perspectival endorsement of the previous source of my belief, but is that really plausible? In the bank case, for example, would it have amounted to a perspectival endorsement that induction is reliable? It seems implausible, and unnecessary, to attribute any such perspective to me, even an implicit one.

doubt our first-order beliefs but that we *should* doubt them—and moreover that only second-order reflective knowledge will adequately resolve our doubts. The problem then is that this normative pressure will presumably apply not just to the object level but to the second-order level as well. For instance, suppose I believe that my neighbor’s dog is in the yard, that I begin to wonder whether the source of my belief (perception) is reliable in these circumstances, and then as a result of studying the relevant cognitive science I come to believe that the source *is* thus reliable. My first-order belief that the dog is in the yard is now, on Sosa’s view, appropriately enlightened—it is now rationally defensible, or so it seems—but what about now my meta-belief that perception is reliable, based on the results of the lab? Can’t I—indeed, if the earlier reasoning was on target, *shouldn’t* I—doubt that too? Isn’t it possible that all of the experiences in the cognitive science lab were hallucinated? Or that (to take one of Sosa’s own recent scenarios) all of my reflective reasoning was performed under the influence of the *Disablex*, a pill that “induces a persistent illusion of coherent empirical reality” (2011: 153)? Based on the Pyrrhonian reasoning, one would think that the same normative pressure to acquire a rational defense against doubt applies just as much on the second-order level as it did on the object level. But then how in the world could I do that?¹⁰

Sosa’s response to this concern is apparently to reject this last step and deny that the same normative pressure is, in fact, felt at the meta-level. As he writes in response to a similar objection by Barry Stroud:

However, it is not easy to understand this [Stroud’s] position. If our perceivers believe (a) that their perception, if reliable, yields them knowledge, and (b) that their perception is reliable, then why are they restricted to affirming only the conditional, a, and not its antecedent, b? Why must they *wonder whether* they understand their relevant knowledge? Indeed, to the extent that they are really convinced of both a and b, it would seem that, far from being logically constrained *to* wondering whether they know, they are, on the contrary, logically constrained *from* so wondering. After all, first, if you are really *certain* that p, then you cannot well consider whether you know it without thinking that you do. Moreover, second, isn’t it incoherent to be convinced that p *and yet* wonder whether p? (2009: 201-202).

The general strategy, then, seems to be to deny that doubt is cogent at the meta-level because of our conviction that sense perception *is* reliable.

¹⁰ There are of course strategies. With Descartes one could argue a priori to the claim that God is no deceiver, hence would not allow it. Or, to take Sosa’s interesting recent strategy, one could argue transcendently that the *Disablex* (at least) is impossible. I do not think those strategies work, but for present purposes it enough to note that even on Sosa’s own view I do not think they are necessary, for reasons I will note in the next paragraph.

Now we are in a puzzling situation, however, because if this move works at the meta-level then it is unclear why it does not work at the object level as well. That is, why not deny there too that doubt about whether the cup is on the table, or the bank is open, or the keys are by the toaster is cogent because we are convinced (let's stipulate) about all of these things? But if this move is unsatisfying at the object level, perhaps because it smacks too much of dogmatism, then it is unclear why it is suddenly satisfying on the meta-level.

In short, if one thinks that something is normatively lacking in a belief unless it is able to respond to doubt, then (as argued earlier) this norm at least in many cases seems satisfiable at the object level, without any need to appeal to a meta-perspective. Alternatively, if we need to turn to the meta-level to resolve doubt, then it is unclear why one does not need to turn to the meta-meta level to resolve further doubts as well.¹¹ But there seems to be no limit here, so in order to uncover the value of reflection it seems best to turn elsewhere.

IV. Intellectual autonomy

We have discussed a few different reasons why Sosa claims that reflective knowledge¹² that *p* is epistemically better than stand-alone animal knowledge that *p*—that only reflective knowledge avoids meta-incoherence, or allows for understanding, or rational defensibility, or the resolution of doubt—and found all of these wanting. Let us now return to the interesting case of the person who neglects to reflect or perhaps even declines to reflect—in the later case, the person who says “No thanks, I’m fine.” Surely *something* seems wrong or lacking about such a response. But what is it, exactly?

In this last section I want to consider what is by my lights Sosa’s most intriguing claim about the value of reflective knowledge, viz., that only by reflecting can we ensure that our beliefs are properly creditable to us, or that we are believing in a way that respects our intellectual autonomy. And previewing a bit, I want to defend two things here. First, that while taking a reflective perspective on our beliefs is indeed importantly connected to autonomy, Sosa misidentifies the point at which autonomy enters the story. It is not by reflecting on the reliability of our faculties that we acquire intellectual autonomy, or that our beliefs become properly creditable to us. Rather, it is by taking our “given” ends as believers to be good or worthwhile that we acquire this autonomy. The second

¹¹ From discussions with Sosa, I now think that his preferred way to solve this problem is by claiming because “ought implies can,” our obligations on this score will quickly peter out. That is, we might be too cognitively limited, after a certain point, to be able to reflect further. It seems to follow for Sosa, however, that a cognitively perfect being such as God would still be required to reflect further, and indefinitely, which (after a few levels of iterations) does not seem like a great result.

¹² In Sosa’s favored sense of reflective knowledge “supported by a structure of reasons.” That qualification should be understood in what follows.

point is that once the connection between reflection and autonomy is properly understood it becomes much less clear whether reflection is properly thought of as an epistemic good or whether it is instead best thought of as a moral or ethical good. Although this classificatory question is in some ways of secondary interest, it is worth making because it helps us to better understand the significance of reflection in our lives. In order to appreciate the difference between our proposal and Sosa's, however, we should first look more closely at the way Sosa connects reflection and autonomy.

We know that for Sosa animal knowledge arises when one's beliefs reliably track how things stand in the world (or, more elaborately, when the accuracy of these beliefs manifests one's ability to track how things stand in the world). But notice that taken in this way one's animal knowledge might seem less than fully human, because not terribly different from the way a good thermometer tracks the ambient temperature, or the way a good supermarket door tracks the presence of nearby shoppers.¹³ Sosa expands on this concern, as well the way in which the appeal to reflective knowledge is supposed to address the concern, in the following passages:

Why the pride of place for reflective knowledge? One answer is to be found in the special bearing of reflective knowledge on the understanding and coherence dear to intellectuals,¹⁴ and on *the intellectual agency that we honor....* No matter how much we value consultation, *we are unwilling to yield our intellectual autonomy*, which requires us to assess the place of consultation in the light of all our other relevant information and recognized desiderata.... In the end reflection has a closer, more finally determinative influence on the beliefs we form, and the deliverances of consultation bear properly only through reflection's sifting and balancing. (2004: 291; italics added)

One question remains pertinent: What could possibly give to reflective knowledge a higher epistemic status than the corresponding unreflective cognitio(n)?... What favors reflective over unreflective knowledge? Reflective acquisition of knowledge is, again, like attaining a prized objective guided by one's own intelligence, information, and deliberation; unreflective acquisition of knowledge is like lucking into some benefit in the dark. The first member of each pair is more admirable, *something that might be ascribed admiringly to the protagonist, as his doing*. And we can after all shape our cognitive practices, individually and collectively,

¹³ As Sosa claims: "Nevertheless, a mere thermometer reaction to one's environment cannot constitute the best human knowledge, regardless of whether that reaction is causally mediated by experience" (Sosa 2009, p. 153).

¹⁴ Here we see another appeal to understanding and coherence, but since I have addressed these earlier I will not return to them again.

enhancing their epistemic virtue, their enabling us to grasp how matters stand. (2009: 142; italics added)

The suggestion is therefore that while in one sense my animal knowledge is ascribable to me because it manifests my own abilities (good eyesight, good memory, etc.), in a fuller or more accurate or perhaps more *important* sense Sosa seems to think it is not my own doing. Why? Because again the beliefs that constitute this knowledge seem to just brutally arise in me, much like the readings of a good thermometer just brutally arise in response to the temperature. My beliefs can only be significantly or importantly ascribable to me— ascribable to me “as my doing”—when I have reflectively endorsed them, which in this context means endorsed them as reliably oriented to the truth.

On this approach reflective knowledge seems to acquire its special value due to its connection to agency and autonomy, and the value of “intellectual autonomy” in particular. It is not entirely clear how having a second-order perspective on one’s first-order reliability is supposed to yield these goods, however. Consider a large supermarket with a variety of doors, all of which are connected to a central system monitoring their performance. When the system determines that a given door is reliable, it allows it to continue functioning; when it determines that a door is unreliable, it shuts it down and calls for repair. If the first-order door reliability was not autonomous to begin with, because a simple brute response, how might it suddenly become autonomous with the addition of second-order monitoring—which, we can suppose, is itself merely a brute response to stimulus, albeit now stimulus regarding how the first-order door systems are performing?

Of course this example is in some ways misleading because with the door there are no beliefs on the scene to begin with—beliefs which might be elevated beyond the animal. But still, the first-order endeavors of the door are nonetheless performances on Sosa’s view, and the question remains as to how brute reactions on the first-order level become performances of a significantly different kind with the addition of a second level brutally responding to stimuli from below. But if this way of connecting reflection and autonomy is unsatisfying, where should we look instead?

The alternative that I would like to sketch instead is I think closer to the earlier Socratic claim that the unexamined life is not worth living, and to the Delphic claim to “Know Thyself.” To see how, suppose for the sake of argument that truth is the natural or default aim of belief and happiness is the natural or default aim of desire. Even though our beliefs and desires might be reliably oriented towards these ends—even though they might track these ends—one might nonetheless worry, along with Sosa, that there is something less than fully human about simply believing and desiring in this way because on this scheme our beliefs would be responding to indications of truth, and our desires to indications

of happiness, in much the same way that a supermarket door would be responding to indications of shoppers.

So what would it take to move us beyond this animal or servo-mechanical level? Plausibly, what is needed is a second-order endorsement of our first-order ends. Even if it turns out that we cannot choose for our beliefs to be oriented towards something other than truth, or for our desires to be oriented towards something other than happiness, we can at least on the reflective level take these ends to be worthy or deserving ends, and we can approve of and try to beneficially shape our various ways of pursuing these ends. As I am imagining it, this sort of reflective approval will not necessarily improve upon the reliability of our beliefs (to say nothing of desires), and hence will not obviously qualify as a kind of “epistemic” improvement, but they will make them more fully our own. Or indeed, they will make them genuinely our own for the first time. It is for this reason, as I suggested earlier, that the sort of reflective perspective just sketched is better thought of as a moral or ethical gain, rather than an epistemic one.¹⁵

Although I am not qualified to say that this is what Socrates had in mind by his claim that the unexamined life is not worth living, it is at least an appealing way of making sense of this provocative idea. On this view, the reason why the unexamined life is not worth living is because it is unfree. It is a life that would be worth living, perhaps, if we were simply confined to our first-order animal perspective, unable to reflect on and endorse (or criticize) our first-order ends. But human beings are not like that, and our dignity consists, arguably, in the fact that we do not merely have beliefs that are naturally responsive to the truth, but we can look down on the reflective level and take this natural responsiveness to be a good thing. We can endorse it or say yes to it.

Of course all of this gives rise to some puzzles [familiar from Frankfurt?]: for example, if the belief at the object level was not really my own doing, because just a kind of brute response, then why think that the belief at the meta-level—wherein I evaluate and perhaps endorse my first-order ends—is somehow more my own doing, or maybe even *for the first time* my own doing? If this belief too is a kind of brute response to ends that I take to be worthy, how can this be the

¹⁵ A side note on the meaning of “epistemic”: what might it mean for one belief, one believing, to be epistemically better than another? I don’t know what else this could mean except that the one belief is better oriented to the truth than the other—in the sense that it is more reliably formed, or more likely to be correct. But is it really true that beliefs guided by a reflective perspective—in either Sosa’s sense or the sense suggested here—are more likely to be correct than beliefs that are not so guided? This does not seem to be the case, at least not across the board. Consider again my belief, via my excellent vision, that my neighbor’s dog is in the yard. Does this belief somehow become more reliable after I come back from the cognitive science lab and learn that my vision is excellent? It does not seem so. But then we should reject the idea that acquiring reflective knowledge necessarily brings about some epistemic benefit. Better instead to explain the value of reflection along other lines, such as attempted here.

source of my freedom or autonomy? As far as I can see, however, this account is the most promising way to defend Sosa's insistence on the importance of reflection, and for the fully human way of living it makes available.¹⁶

Works Cited

- Aaron P. Blaisdell et al. 2006. "Causal Reasoning in Rats." *Science* 311: 1020-22.
- Kornblith, Hilary. 2004. "Sosa on Human and Animal Knowledge." In *Ernest Sosa and His Critics*. Ed. John Greco. Malden, MA: Blackwell.
- Sosa, Ernest. 2004. "Replies." In *Ernest Sosa and His Critics*. Ed. John Greco. Malden, MA: Blackwell.
- . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. New York: Oxford University Press.
- . 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. New York: Oxford University Press.
- . 2011. *Knowing Full Well*. Princeton: Princeton University Press.
- Taylor AH, Hunt GR, Medina FS, Gray RD. 2009. "Do New Caledonian crows solve physical problems through causal reasoning?" *Proc. Biol. Sci.* 276: 247-54.
- Tiberius, Valerie. 2010. *The Reflective Life: Living Wisely With Our Limits*. New York: Oxford University Press.

¹⁶ Thanks to Allan Hazlett, Ernest Sosa, Miguel Fernandez, and two anonymous referees for helpful comments on earlier versions of this paper.