

---

## Critical Notice

ERNEST SOSA. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press 2007. Pp. xiii + 149.

Sosa's influence on contemporary epistemology is hard to overstate. The position that he articulates in *A Virtue Epistemology* is a magnificent contribution to the field, one that synthesizes with great effectiveness a number of main themes, and that provides further elaboration and development of the position that he defends in *Knowledge in Perspective*.

I will begin by surveying some main themes in *A Virtue Epistemology*: the view of animal and reflective knowledge that Sosa develops, his two responses to skepticism, and his credit theory of knowledge. Then I will try to determine in more detail what animal and reflective knowledge are, for Sosa, and show some problems that arise from his conception of them.

At the time of writing, Sosa's book had already been out for a year and a half, and a large amount of critical material has already appeared in print, along with replies from Sosa. I will take the liberty of making use of some of this material in this notice, since much that is difficult in the book is clarified in Sosa's responses to his critics.

### I Apt belief and reflective knowledge (a first pass)

Performances can be judged on what Sosa calls the AAA model: they can be *accurate*, or successful; *adroit*, or skilful; and *apt*, accurate because adroit, i.e., successful because they were performed skilfully. Beliefs,

he argues, are long-sustained performances, and can exhibit the AAA structure as well.

Adroitness, or whether a belief arose from a competence operating in the right conditions (whether or not it was successful), is a form of justification. For instance, Sosa argues that the best way to understand the justificatory power of rational intuitions is to see them as the deliverances of epistemic competences. This model, he maintains, is superior to treating intuitions on the model of perception, or as grasping intellectual facts.

Apt belief is belief that is true because it manifests an epistemic competence exercised in appropriate conditions. When one's belief is apt, it would not likely be false as long as the conditions are appropriate and one's competence is preserved; and this highly relativized safety arises because of, or is attributable to, one's competence. The conditions are appropriate when either (a) the subject is guided by them, as when an archer corrects for the wind when shooting or a believer forms a visual belief because the light seems good, or (b) they are constitutive of the normal conditions for exercise of the competence, as when an archer shoots on a calm day without any regard for whether the wind is good or a believer forms a visual belief in good light without any regard for whether the light is good.

Sosa has previously advocated<sup>1</sup> that knowledge must be safe: to know that  $p$ , one's belief that  $p$  must not easily be false, i.e., must not be false in nearby possible worlds. Aptness does not entail safety, because one's competence or the appropriateness of the conditions might be counterfactually fragile. If one's ability might easily have been impaired, or one might easily have been in unfavourable conditions, one's belief will not be safe. The move from safety to aptness is parallel in many ways to Nozick's relativization of truth-tracking to the process type that actually formed the belief. We need first to take account of the fact that one's belief can be unsafe because one was lucky to have the evidence one had, which gets us basis-relative safety, or safety relative to the grounds for the belief. We must also, Sosa thinks, take account of the fact that one might be lucky to be able to exhibit the competence that one did, or might be lucky that one was able to exhibit that competence in the right conditions. Consider a nervous student in an oral exam who forgets his nervousness just long enough to be able to solve a complicated problem that he can only solve when relaxed. It seems that the student knows the solution to the problem, even though he might easily have been too nervous to compute it correctly, and gotten a wrong answer.

---

1 E.g., in Sosa 1996.

Apt belief is animal knowledge, the lowest or least valuable sort of knowledge. A better, more valuable state is reflective knowledge (RK), which is 'apt belief aptly noted' (32). For RK, one must ably recognize the aptness of one's beliefs, or be in a position to know or perhaps rationally defend them. Sosa is rather vague on just what RK is; we'll first look at the use he makes of it in the book and then try to get a clearer sense of it.

## II The imagination model of dreaming

The guiding theme of the book is a response to skeptical scenarios, the broad outline of which seems to be this. Sosa rejects brains-in-a-vat and evil-demon scenarios on the grounds that they are too distant to be relevant possibilities. They depend on imposing an illegitimate requirement of sensitivity (that one not believe that  $p$  were  $p$  false) on knowledge, but do not challenge safety or aptness. Thus, they do not challenge our knowledge.

Dream scenarios, however, are 'too close for comfort' (2). Since we dream nightly, it's hard to rule out as too distant the possibility that we might be dreaming right now. However, in the sort of dreams that are close possibilities, we would lack the epistemic competences that we have when we are awake. Thus no matter how close the possibility that we are dreaming is, it does not undermine our knowledge (animal or reflective). Sosa presents two versions of this general strategy.

The first version is the 'imagination model' of dreaming. On the standard conception of dreaming, when we dream we have false beliefs; for instance, if I dream I am on the moon, I believe falsely that I am on the moon. Sosa proposes, rather, that to dream one is on the moon is to imagine oneself on the moon.

On the imagination model, the belief that one is not dreaming is safe regardless of how close the dreaming possibility is, because if one were dreaming, one would merely imagine oneself to be awake, rather than believing oneself to be awake. Thus the proposition 'I am awake' is self-intimating, rather like the cogito; if believed, it must be true.

Of course, the cogito has an additional advantage over the proposition 'I am awake'. We can be directly aware of the fact that we are thinking, and thus directly aware that we exist. Sosa grants that the content of a dream-state might be subjectively indistinguishable from that of a waking state. Although we cannot believe ourselves to be awake unless we are awake, we also cannot tell just by reflection whether we actually believe ourselves to be awake or are merely imagining ourselves to be awake. Without some way of distinguishing those two possibilities

internally, it seems we cannot have reflective knowledge that we are awake.

If one reflects on the choices available, Sosa thinks, one can see that believing oneself to be awake is the epistemically rational act. We can lay out a decision table as follows:

	<i>Try to believe</i> <sup>2</sup>	<i>Try to suspend judgment</i>	<i>Try to disbelieve</i>
<i>Awake</i>	Believe oneself awake; true belief	Suspend; no error	Disbelieve; false belief
<i>Asleep</i>	Imagine oneself awake; no error	Imagine suspending; no error	Imagine oneself asleep; no error

Imagining oneself awake is, for Sosa, no more of an error than day-dreaming that one is in Bimini. As long as  $v(\text{true belief}) > v(\text{no error}) \geq v(\text{false belief})$ , trying to believe dominates.

### III A second response to dream skepticism

The imagination model requires an unorthodox account of dreaming, but is compatible with a safety requirement for knowledge; the second version of Sosa's response to dream skepticism is compatible with the orthodox view of dreams, but requires replacing safety with aptness. The normal conditions for the exercise of our epistemic competences include being awake. If one is dreaming, one is not in appropriate conditions for the exercise of one's competences. Hence, the possibility that one might be dreaming no more undermines the aptness of one's beliefs than the possibility that an archer might be poisoned and rendered incapable of accurate shooting undermines the skilfulness of his shots *now*.

Mere apt belief is small consolation, however. Consider a subject who is perceiving a red surface in good light; but it is a kaleidoscope surface, and both it and the ambient light are controlled by a joker who might

---

2 I mean 'trying to believe' to be an intellectual act that leads to believing that p when awake and imagining that p when asleep. Sosa maintains that when one is dreaming, one cannot actually try to believe any more than one can actually believe (2009a: 429).

just as easily have presented the subject with a white surface under red light. This example is closely parallel to the better-known scenario in which Barney sees a real barn in an area with mostly barn façades. In both cases, the subjects satisfy the requirements for apt belief. Their competences for colour-attribution and barn-recognition are operating in appropriate circumstances, and would not easily have issued false beliefs as long as the circumstances were appropriate. The circumstances might easily have been inappropriate had the joker presented a white surface under red light, or had Barney happened to look at a fake barn, but that only undermines safety, not aptness.<sup>3</sup> We are strongly inclined to deny knowledge in both cases, so this conclusion is problematic.

Sosa accounts for our intuition by proposing that while these subjects have apt belief, they lack reflective knowledge. They do not aptly believe that they know, or (perhaps better) are not in a position to aptly believe that they know, were they to reflect on the belief in question. To know that they know, the subjects must have a metacompetence that aptly indicates that their belief issued from a competence under appropriate conditions. In cases like these, our metacompetence takes the conditions to be favourable absent any sign to the contrary. Since the conditions could easily have been unfavourable without the subject's having any sign that they were, this default competence is right only by luck; the correctness of the subject's metabelief is not or would not be attributable to an ability to detect when the circumstances are unfavourable. Thus the subjects do not know that they know.

Several points on this response are worth highlighting before continuing. First, had the joker presented a white surface under red light, the subject's colour-recognition ability would no longer be operating under normal conditions, but her metacompetence would still be. Thus the normal conditions for metacompetences are broader than those for the object-level competences they oversee. This is just as it should be. Surely, you can't have an ability to tell when an object-level competence is impaired or in abnormal conditions that is *itself* impaired or in abnormal conditions whenever the object-level competence is. This extra breadth might arise because the normal conditions for metacompetences are individuated largely by one's internal states — by features of one's experience, such as the apparent light level, rather than by fea-

---

3 On many accounts, however, the fact that the conditions might easily have been different means that the success is not attributable to the *competence*, but rather to one's good fortune of having been in the right conditions. See, e.g., Zagzebski 1996: 293-9, Greco 2003, Riggs 2007. Sosa's account of a competence is weaker than most virtue epistemologists', a point I'll return to below.

tures of the external world, such as the ambient light. Thus when the joker changes the ambient light but not the visual experience, conditions are abnormal for the object-level competence but not for the meta-competence.<sup>4</sup>

Moreover, this broadening of normal conditions means that reflective knowledge comes much closer than apt belief to being safe. The broader we make the normal conditions for a competence, and the more broadly we individuate that competence itself, the broader the range of possible worlds across which it is possessed, and the closer its aptness comes to basis-relative safety.

Finally, note that in accounting for our intuitions that these subjects do not know by granting them animal knowledge but denying them reflective knowledge, Sosa seems to identify our intuitions about 'knowledge' with the latter rather than the former. I'll take this up in more detail below. For present purposes, what matters is that merely showing our perceptual beliefs apt despite the possibility of dreams is insufficiently satisfying. We need also show that we have reflective knowledge, so that ordinary perceptual beliefs do not end up having the same status as those in kaleidoscope/fake barn cases.

Dreams, Sosa notes, are characteristically unlike waking life. They have what Austin calls a 'dream-like quality' to them, and they lack the coherence of reality. Granted, when we are dreaming, we often cannot tell that we are not awake. But this, Sosa proposes, is because our capacities for telling whether we are in appropriate conditions for trusting our perceptual faculties are *themselves* impaired. Hence, our inability to realize that our perceptual faculties are misleading us when we are dreaming does not entail that we lack the competence to know that we know when we are awake.

Although he doesn't present it this way, it seems that this response is at best a lemma. In the dreams we dream regularly, and that we can claim are 'too close for comfort,' our cognitive and metacognitive abilities are impaired or operating in abnormal conditions, and thus the aptness of our waking beliefs and metabeliefs is not challenged. It seems possible that we could be in a dream that is subjectively indistinguish-

---

4 In a reply to his commentators, Sosa says the appropriate conditions for the meta-competence are 'that one be appropriately alert, attentive, and possessed of perceptual and cognitive faculties that would detect tell-tale signs (*internal* signs, since we are granting this much internalism for the sake of argument) contrary to the satisfaction of those first-order requirements' (2009b: 146, emphasis in original). Unfortunately, the qualification that internalism is granted 'for the sake of argument' leaves it unclear just to what extent the conditions for metacompetences are internally determined.

able from reality even when our metacompetences are *not* impaired. But the actual phenomenon of dreaming gives us no reason to think *this* possibility is close enough to endanger our knowledge.

#### IV A Cartesian response to Pyrrhonian skepticism

In the final chapter of the book, Sosa proposes a response to the Cartesian Circle.<sup>5</sup> The mistake often attributed to Descartes is that he held: (a) that propositions rationally intuited or deduced from rationally intuited premises are known for certain; (b) that nothing is known for certain until God's existence can be proven; and (c) that God's existence must be proven through appropriate reasoning. These three are inconsistent, of course.

Sosa's resolution of the inconsistency is to propose that Descartes' distinction between *cognitio* and *scientia* (found in the *Objections and Replies*) corresponds, approximately, to his own distinction between apt belief and reflective knowledge. Intuition and deduction from the intuitive give us *cognitio*. Once we have *cognitio* of God's existence, however, we see that the various things we have *cognitio* of are not just true by luck. That lets us bootstrap into *scientia*, which is a higher, more valuable epistemic state. The process is circular if we suppose that we must have *scientia* of the premises from which we must prove God's existence before we can have *scientia* of anything at all. But there is no circularity if *scientia* must only be preceded by the lesser state of *cognitio*.

I'll leave the question of whether Descartes held this view to others. What's more interesting for my purposes here is what it tells us about Sosa's view. Apt belief provides foundational knowledge; it seems that if a belief is apt, one needs no further reasons for holding that belief, drawing implications from it, etc. At this level, the response to skepticism is purely externalist, since we merely reject the skeptic's demand that reasons for believing accompany apt belief.

However, apt belief is also deeply unsatisfying; it is an epistemically valuable state, but not particularly so. 'Knowing full well', as Sosa calls it in this chapter, requires some sort of awareness of why one's belief is justified. But since we are entitled to hold apt beliefs, we can make use of them in the course of acquiring a perspective that gives us this awareness of why we are justified. That elevates us into a higher epistemic state, reflective knowledge. Once we have reflective knowledge, we can aptly identify reasons for our apt beliefs, and then we have a full response to the skeptic.

---

5 Much of the material in this chapter is discussed in more detail in Sosa 1997.

## V Epistemic normativity

Aptness also plays the central role in Sosa's account of why knowledge is more valuable than mere true belief. He holds that the fundamental value of epistemic appraisal is truth. Epistemic norms concern whether one's beliefs are true, whether they were formed in a way likely to render them true, and so on. The norms are entirely neutral, however, on the broader question of what beliefs one should have, or even whether one should have beliefs at all, as long as the beliefs that one does have are true. The more general questions, Sosa thinks, are a matter for a broader study he calls 'intellectual ethics,' which he refrains from speculating on.<sup>6</sup>

If truth is the fundamental epistemic value, it is hard to see why reliably or safely produced true belief should be better than lucky true belief. One maximizes one's chance of getting true beliefs by using reliable or safe methods, but it is hard to see why the end state of reliably produced true belief should be any better than the end state of a lucky guess.

To resolve this problem, Sosa endorses the credit theory of knowledge. As I noted above, he takes beliefs to be sustained performances. In general, a successful performance is better by virtue of succeeding through ability rather than just luck. A shot that hits the target because of the skill of the archer is better than one that hits the target just because of a lucky gust of wind. Thus the exercise of an intellectual competence is not just of instrumental value because it makes one more likely to get true beliefs; the whole performance of believing correctly because of one's own skill is more valuable than that of believing correctly out of luck. This response to the value problem gives a further reason for requiring aptness rather than safety for knowledge. The greater value of apt belief over luckily true belief can be explained more readily than the greater value of safe belief.

Lackey (2007) has argued that the credit theory can't account for testimonial knowledge. Suppose Morris asks the first passerby he sees how to get to the Sears Tower, and acquires a true belief as a result. It seems that Morris's own contribution to the truth of his belief is too small for him to get credit for it; rather, it would seem that the credit should go to the informant.

---

6 Although the assumption that he needs in his defense of reflective knowledge on the imagination model of dreaming — that truly believing oneself awake is better than not erring — is a substantive claim in intellectual ethics.

Sosa's response to this is to argue that many of our beliefs are generated by social competences that are partially realized in many different people. The competence that gave rise to a particular goal in soccer might have involved the orchestrated contributions of many different players, each of whom receive partial credit for the outcome. In like manner, an epistemic competence might involve a research team doing a study, a reliable journalist who reads the study and summarizes it in the *Economist*, and my reading that article, understanding it, and coming to have a true belief as a result. Each of the persons involved receive some degree of credit for the truth of my ensuing belief. Obviously, most prominent among these will be the team who originally made the discovery. But I still receive enough credit for the final outcome to have knowledge.<sup>7</sup>

## VI Apt belief and reflective knowledge, second pass

Let us return to the problem I set aside earlier: what exactly is the distinction between apt belief (AB) and reflective knowledge (RK)? And where does ordinary knowledge (OK), the sense in which we typically use the word 'knows,' fit in?

Any attempt at an answer to the latter question must begin with a caveat. Sosa has for some time been quite concerned with reconciling, as much as possible, the internalist and externalist camps in epistemology. The AB-RK distinction is intended in part to capture different standards for knowledge that different schools of thought might endorse. Externalists are happy with something like AB, while internalists require something more like RK, and the Pyrrhonian skeptic expects a very demanding sort of RK. Thus, if we ask Sosa, his answer to the second question would almost certainly be that there isn't a single sense in which *all* of us use the verb 'knows'.

Setting skepticism aside, we can note that internalists and externalists do agree on a very wide range of cases. While they tend to differ on, e.g., reliable clairvoyants, I don't think there is much disagreement that Gettierized beliefs are not knowledge, while clear cases of perception, rational intuition, deduction, and the like are. It seems that the standards for knowledge can also depend on how high the stakes are;<sup>8</sup> let us set aside such cases as well. So let us understand OK to be what

---

7 See Lackey 2009 for a response to Sosa's strategy here.

8 See Hawthorne 2004.

is deployed in all the cases on which there is reasonably widespread agreement, and little apparent influence from practical concerns.

*A Virtue Epistemology* suggests three answers to those two questions:

1. RK is OK, while AB is 'brute cognition' (2009a: 431), or knowledge in a 'metaphorical' (432) sense. AB isn't really knowledge, but it is still the case that we need no additional reasons for holding it. Thus AB provides externalist foundations from which we can bootstrap ourselves up to knowing.
2. AB is knowledge; RK is knowing that you know (KK). Ordinary usage doesn't make much of a distinction (if at all) between knowing and knowing that one knows; moreover, our intuitions (or at least my intuitions) about what beliefs fall on what side of the distinction aren't very clear. Thus, on this possibility we might also say that OK is closer to KK than to knowledge, that knowledge without KK is something of an externalist term of art.
3. AB is OK; RK is a superior epistemic state, one that is better than mere knowledge. RK is reflective and rationally defensible. It is our aim in most of our intellectual endeavours to attain RK as much as we can. We might even say that RK is closer to 'understanding' than to 'knowledge.'<sup>9</sup>

These are not mutually exclusive, of course. In *A Virtue Epistemology* one can find textual evidence for all three of these accounts, and little reason to favour one over the others.

Elsewhere, Sosa (2004: 294-301, 2009a: 431-2) has given a more detailed account of AB and RK,<sup>10</sup> and it seems best to turn to those sources to understand the position Sosa takes in the book. On this fuller account, there is a spectrum from apt belief to reflective knowledge. At the bottom level, we have true beliefs formed in appropriate conditions. The believer may not, however, display any sensitivity to whether the conditions are normal for the operation of that competence. Thus a deer that runs from a nearby predator believes aptly that it is in danger, even though it would have formed the same belief were it approached by a vegan. This is just knowledge in a metaphorical sense, the sense in which we say that thermometers 'know' the temperature and automatic doors 'know' when someone wants to come in.

---

9 This is a point that Grimm 2001 proposes.

10 To which he adverts in *A Virtue Epistemology*, saying '[a]ny full account would have to register how [AK and RK] are matters of degree' (32).

Humans and more intelligent animals, however, can acquire various amounts of sensitivity to the appropriate conditions for their competences. They can learn how to adjust their belief-formation to make it more reliable. They can adapt their belief-formation when their competences degrade or the environment changes. They can acquire a sensitivity to the conditions in which their faculties are reliable, and become more willing to accept those deliverances when trustworthy and less willing to accept them when they are not. This flexibility, Sosa maintains, requires some sort of awareness of the sources of one's beliefs and the conditions under which they successfully operate. That is, it requires metacognition, the capacity to monitor and regulate one's own cognitive processes.<sup>11</sup> This awareness, however, can be entirely implicit; manifested not in explicit beliefs about one's capacities, but just in one's ability to effectively and adaptively regulate belief-formation. As this awareness becomes more explicit one moves up the scale of reflectiveness. As we approach the higher levels, we find better and better capacities to identify explicit reasons for holding certain beliefs. At the highest level, for each belief we have a fully articulable and rational account of why that belief is justified.

Let's identify, with a certain amount of arbitrariness, two particularly important regions on this scale above AB. Let 'implicit reflective knowledge' (IRK) be the region of the scale in which one has a reasonably rich and sensitive implicit awareness of one's faculties, manifested in reasonably effective capacities for metacognitive control, but at most crude and limited explicit meta-beliefs. Let 'explicit reflective knowledge' (ERK) be the region around the top of the scale, in which one has an explicit awareness of one's competences that provides articulable defenses for one's beliefs.

Ordinary human knowers have some metacognitive capacities, and so some sort of awareness of the reasons for which their beliefs are apt. In ordinary unreflective subjects, however, we may suppose that this is mostly along the lines of an implicit awareness. Our capacities, for instance, to evaluate information retrieved from memory and decide whether to accept it go far beyond our typical capacities to *articulate* reasons why we should trust our memories. On these grounds, it seems, we should identify OK with IRK.<sup>12</sup>

---

11 See Carruthers 2008, however, for an argument that capacities to adapt one's faculties and calibrate strength of belief to reliability do not necessarily require any sort of awareness of one's cognition at all.

12 Which is effectively to endorse Henderson and Horgan's 2007 requirement of 'modulational control' for knowledge.

IRK is a bit like knowing that one knows, in that one has meta-competences for judging whether, e.g., the conditions for one's visual perception are appropriate. At this level, however, these metacompetences are primarily manifested in intelligent regulation of object-level competences rather than issuing in explicit beliefs. For instance, one's metacompetence to detect when the conditions for visual perception are normal might manifest itself in an ability to refrain from forming visual beliefs when the conditions are abnormal, rather than in an explicit belief as to the favourability of the present conditions. It might be the case, however, that this awareness is sufficiently accessible that a suitably attentive and reflective believer could come to know that she knows on the basis of the awareness she already has. Thus, it seems, we can still run Sosa's analysis of the dream and kaleidoscope/fake barn scenarios with only minor changes.

In a reply to his commentators, Sosa does seem to endorse the identification of OK with IRK; he says that 'apt belief aptly noted...is particularly associated with rational animals of our order of sophistication' (2009b: 141), while nonetheless distinguishing this state from reflective knowledge '*of a high order*' (141, e.i.o.).

Note that IRK is much closer to basis-relative safety than mere aptness. In IRK, one has at least some awareness of the conditions that are favourable for the operation of the faculty, and the conditions under which the faculty is impaired, and this implicit awareness is manifested by an ability to refrain from trusting the faculty's deliverances under those conditions. The implicit awareness will tend to reduce the likelihood of false beliefs formed in nearby worlds with unfavourable conditions or impaired faculties, and thereby increase the safety of those beliefs. Reflectively informed aptness may not be identical to safety, but it is at least much closer than uninformed aptness.

Apt beliefs are still sufficiently well formed to be held without further reasons, but AB is a lesser state, not sufficient for knowledge. ERK is a higher state. We often apply the word 'knowledge' to cases that are not ERK (though the skeptic perhaps tries to prevent us from doing so) but ERK is the more valuable state, and the most worthy epistemic aim. This is the picture that I will address in what follows.

## VII Problems with apt belief

The first problem for the three-way distinction between AB, IRK, and ERK is that AB seems too weak to do very much work. One wonders if a proper response to the skeptic can begin with foundations that are reliable, but not actually known. Sosa is surely right that we can only avoid vicious circularity by requiring a lesser epistemic status for our

starting-points than what we hope to eventually achieve. But if the starting-points are not *known*, not in the ordinary sense of the word, we seem to be dodging the skeptic's challenge rather than answering it.

Now perhaps that is the right strategy to take with the skeptic. Perhaps we should circumscribe the skeptic's right to issue challenges, and one could probably make a case for denying the skeptic a right to challenge apt beliefs without positive reasons for thinking them false. I think it is clear, though, that this is not Sosa's preferred approach. At levels of reflectiveness intermediate between IRK and full-blown ERK, we might be able to answer skeptical challenges to some beliefs, while to others we can only mutely reject the legitimacy of doubt. (Think of a scientist who can give a detailed account of why her experimental results support her conclusion, but can only dismiss as silly the question why she believes her results didn't just occur in a dream.) Full-blown ERK should, however, allow us to articulate defensible reasons for discounting skeptical challenges in this way. And such an articulation seems much more defensible if it can claim that the starting-points from which we bootstrapped to ERK were things we *knew*, rather than just brutally cognized.

Setting that worry aside, we are supposed to receive credit for apt belief. On Sosa's analysis, subjects in the kaleidoscope/fake barns cases have AB but not IRK. Denying them IRK is supposed to accommodate our disinclination to attribute knowledge to them. AB, however, is supposed to be sufficient for credit. The result is highly unstable; if we are unwilling to say that those subjects know, it seems we should be unwilling to grant them credit for having their true beliefs.

Sosa invokes the credit theory in order to explain why knowledge is better than mere true belief. If AB is enough for credit, however, we have a new value problem: why is ordinary human knowledge, i.e., IRK, better than creditworthy true belief?

Finally, Sosa at times says that apt belief is only 'metaphorically' knowledge. That would suggest that it should line up with cases of metaphorical credit. Consider the first few halting steps taken by a very young child. We are inclined to call that an achievement, to credit the child with having walked for the first time. But the very first performance of walking almost certainly doesn't succeed out of ability. The child lacks an ability to walk just yet; this is just the first instance in which he was lucky enough not to fall over for a few steps. It is, nonetheless, creditworthy in a sort of metaphorical sense, or a very weak sense made appropriate by the child's limited abilities and his parents' desire to laud his development.

Perhaps Sosa thinks that credit is a sliding scale like knowledge; as one becomes more reflective, one becomes more competent, and receives more credit for truth. The preceding considerations show, however, that

we need to move up the scale past AB before we find sufficient credit for knowledge. Sufficient credit for knowledge requires that one have abilities in a stronger sense than is necessary for apt belief.

It seems then that epistemic competences or virtues should be the sorts of abilities that give rise to IRK, rather than those that give rise to merely apt belief. Epistemic competences must exhibit a certain amount of flexibility and sensitivity to whether the conditions for their operation are appropriate and whether their capacities for success are impaired. Intellectual virtues must be guided by an awareness, however implicit, of their own reliability in the circumstances. This awareness, of course, need not constitute knowledge. If it is an implicit awareness, it need only allow the competence to be appropriately flexible and adaptable; it need not be a belief, much less a justified one. On such a strengthened conception of ability, then, apt belief would line up with knowledge in the usual sense of the word.<sup>13</sup>

### VIII Problems with explicit reflective knowledge

The scale proceeding from AB through IRK to ERK is seemingly meant to be linear, but there are two separable dimensions on which one's understanding of one's faculties can increase. On the one hand, one's awareness of one's faculties and the conditions under which they are reliable can improve, as manifested in one's ability to selectively trust and apply them. On the other hand, the explicitness of this awareness can increase, as one acquires more metabeliefs about one's faculties and stronger rational defenses of the beliefs acquired from them.

These two do not necessarily go together. One's implicit awareness of one's competences can improve without its becoming more explicit. One might have an uncanny ability to apply a competence only when it is highly likely to yield true beliefs; but the awareness that permits this might be entirely implicit, so that one has no capacity at all to appeal to it when defending one's beliefs. Consider a chicken sexer whose confidence in her judgments issued under different circumstances unerringly tracks the actual probability of her being right in those circumstances, but who has no understanding of how her confidence is

---

13 One might worry here that the more we require for credit, the harder it becomes to explain why credit is due in Lackey's testimonial cases. Given Sosa's treatment of the kaleidoscope/fake barns cases, to satisfy our intuitions about testimony he needs to establish that subjects with testimonial knowledge have IRK, not just apt belief (cp. Lackey 2009). Hence requiring IRK for credit won't make that job any harder.

so well calibrated. This is generally the case for abilities: it is one thing to be able to make shots from the three-point line under certain conditions; something further to be able to tell when one should attempt a three-point shot; and something distinct from *that* to be in a position to articulate why one attempted a particular three-point shot.

So take some characteristic source of error — a decisionmaking bias, say, or a set of conditions in which one of our competences is characteristically wrong. Suppose A has implicit beliefs that coarsely indicate the source of error. These implicit beliefs reduce the probability that A will make the error, although it still happens sometimes. Suppose B has the same beliefs as A, but hers are explicit. She knows approximately what the error is and what to do to prevent it, and her object-level performance is the same as A's. Finally, suppose C has very detailed, fine-grained implicit beliefs about the error. These enable him to avoid it in nearly all cases, but they are completely implicit; he has no more explicit awareness of the error than A does, although his object-level performance is vastly better.

B and C are both more reflective than A. It's not clear, however, how we can compare the former two. Their metacompetences — one coarse but explicit, the other fine-grained but implicit — just seem to lie on different dimensions.

A further problem is that one's metabeliefs might not actually provide guidance in belief-formation. Conscious deliberation and recalling detailed or complex information are both slow and resource-intensive, placing great demands on attention and working memory. Suppose a subject had detailed meta-beliefs about why he thinks that conditions C are conducive to the formation of trustworthy visual beliefs, such that if challenged, he can give a detailed and eminently rational defense of why visual beliefs formed in conditions C should be accepted. It is hard to imagine that the subject would actually consult those detailed meta-beliefs on the question of whether to accept a certain visual belief. Cognizers like us can't go about forming all our beliefs this way. Besides the difficulty and slowness of explicit reasoning, our explicit metabeliefs often leave out a great deal of data that we can marshal when forming beliefs. Our capacity to recognize faces, for instance, makes use of far more visual information than we can describe.<sup>14</sup>

---

14 In fact, it seems that explicit metabeliefs could actually interfere with our meta-cognitive capacities by leading us to ignore inarticulable data in belief-formation. Subjects who are asked to describe a face after seeing it are less able to recognize that face later on. It appears that the memory of the verbal description can interfere with the more detailed visual memory, leaving the subjects with less information to use in recognition. See Schooler & Engstler-Schooler 1990.

What is far more psychologically plausible is that someone with detailed meta-beliefs has explicit beliefs reflecting or summarizing their implicit awareness, but only the implicit awareness is causally operative in belief-formation. Suppose, for instance, I see what looks like an elm but refrain from believing that it is an elm because of an implicit sense that I might have looked at it too quickly. This might be turned into ERK in two ways. I might have refrained from forming the belief *because* I had the metabelief; that is, I formed the metabelief and then on that basis rejected the object-level belief. Alternatively, I might believe that I saw the tree too quickly to form a secure judgment, but this meta-

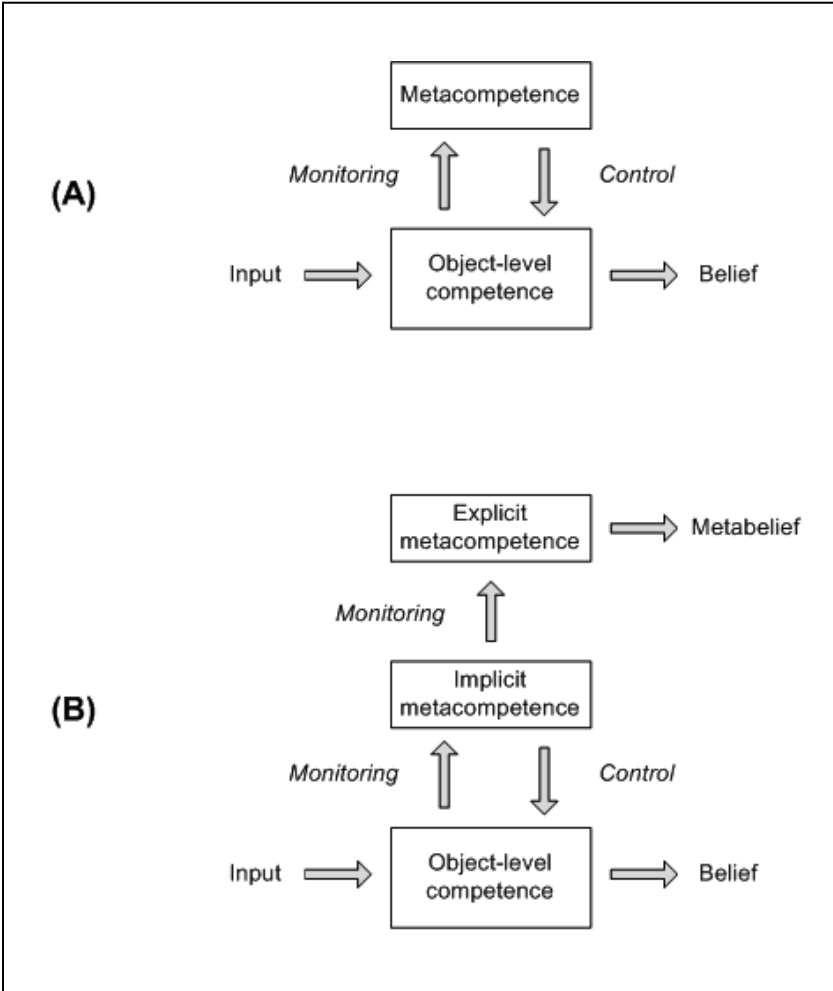


Figure 1. Two roles for metabeliefs.

belief was not causally operative; rather, it was a report or observation of the causal influence that my implicit awareness had.

It may help to diagram the two conceptions (Figure 1). In (a), the subject has a metacompetence (involving an implicit or explicit awareness) that is involved in belief-formation by virtue of monitoring and controlling the object-level process. In (b), the metacompetence is assumed to be implicit; an explicit metacompetence observes its behaviour and generates on that basis a metabelief. We assume here, however, that at least in normal circumstances the explicit metacompetence is *not* itself causally operative in belief-formation.

It makes sense to see (a) as a single performance. In (a), the metacompetence is intimately involved in the formation of the belief. Moreover, the extra value that arises from the competence being guided by a metacompetence is closely linked to the value of truth. It enhances the security of the belief by making it safer, less likely to have been formed in inauspicious conditions or when the object-level competence is impaired. Thus it is not hard to see why it is better for a performance to be intelligently guided than not, just as a three-point shot made by a player who knows when he should make an attempt is a better performance than one by a player who merely happened to be in the right conditions when he tried.

In (b), it seems we have *two* performances: one issuing in a belief, and a second one that observes the first and issues in a metabelief. But then we have the following problem: why should the status of the *object-level* belief be any better just because its generation was aptly observed? The object-level belief is no more likely to be true, no safer, by virtue of being aptly observed; why, then, should we think that it is more valuable?

It seems, then, that in order to hold up ERK as the highest state, Sosa has to introduce a second epistemic value, that of being in a position to rationally defend one's beliefs. Unlike effective regulation, this seems irreducible to the epistemic value of believing correctly or ably generating correct beliefs.<sup>15</sup> What we need is an account of why true metabeliefs are valuable, even though it's not generally the case (for Sosa) that true beliefs are better than no beliefs.

In *Knowledge and Perspective*, Sosa proposes that beliefs can be evaluated by two standards: the externalist standard of safety, and the internalist standard of perspectival coherence. The highest epistemic good, reflective knowledge, satisfies both standards. One problem with that

---

15 Note that the same problem will arise for, e.g., Alston's (2005) attempt to derive the value of conscious access to the grounds of one's beliefs from truth-conduciveness.

view is that the two are independently satisfiable, and nothing ties them together.<sup>16</sup> The modified view that Sosa presents in *A Virtue Epistemology*, with the conception of the implicit perspective, the idea that knowledge is credit for attaining true belief through ability, and the idea that reflective knowledge is 'apt belief aptly noted', does not do enough to unify the two standards. Internalist desiderata and externalist truth-conduciveness are still at loggerheads, even within Sosa's theory.

## XI Conclusion

In conclusion, there are two chief points that I would urge on Sosa:

- 1) Merely apt belief is too weak to play the role Sosa assigns to it. Ordinary human knowledge requires at least an implicit awareness of one's competences and their trustworthiness in different circumstances, manifested in a sort of informed flexibility in the way one treats their deliverances. Since that sort of capacity seems necessary for knowledge, the credit theory seems to require that it be necessary for having an epistemic competence or virtue.
- 2) The ability to use a faculty intelligently, in a way that is sensitive to the conditions under which it safely yields true beliefs, and the ability to articulate or defend the deliverances of that faculty, are distinct. Hence, Sosa needs an account of why rationally defensible apt belief is better than merely apt belief, or why it is valuable to have true beliefs about the grounds of one's beliefs.<sup>17</sup>

*Received: September 2009*

CHRISTOPHER LEPOCK  
University of Toronto  
Toronto, ON M5R 2M8  
CANADA

---

16 This is argued in Lepock 2006 and Bernecker 2006.

17 I thank Jennifer Nagel and Ernest Sosa for helpful comments and suggestions.

## References

- Alston, W. P. 2005. *Beyond 'Justification': Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Bernecker, S. 2006. 'Prospects for Epistemic Compatibilism'. *Philosophical Studies* **130**: 81-104.
- Carruthers, P. 2008. 'Meta-Cognition in Animals: A Skeptical Look.' *Mind and Language* **23**: 58-89.
- Greco, J. 2003. 'Knowledge as Credit for True Belief.' In *Intellectual Virtue: Perspectives from Ethics and Epistemology*. M. DePaul and L. Zagzebski, eds. Oxford: Oxford University Press.
- Grimm, S. R. 2001. 'Ernest Sosa, Knowledge, and Understanding.' *Philosophical Studies* **106**: 171-91.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Henderson, D., and T. Horgan. 2007. 'Some Ins and Outs of Transglobal Reliabilism.' In *Internalism and Externalism in Semantics and Epistemology*. S. Goldberg, ed. Oxford: Oxford University Press.
- Lackey, J. 2007. 'Why We Don't Deserve Credit for Everything We Know.' *Synthese* **158**: 345-61.
- \_\_\_\_\_. 2009. 'Knowledge and Credit.' *Philosophical Studies* **142**: 27-42.
- Lepock, C. 2006. 'Adaptability and Perspective.' *Philosophical Studies* **129**: 377-91.
- Riggs, W.D. 2007. 'Why epistemologists are so down on their luck.' *Synthese* **158**: 329-44.
- Schooler, J.W., and T.Y. Engstler-Schooler. 1990. 'Verbal Overshadowing of Visual Memories: Some Things are Better Left Unsaid.' *Cognitive Psychology* **22**: 36-71.
- Sosa, E. 1996. 'Postscript to "Proper Functionalism and Virtue Epistemology".' In *Warrant in Contemporary Epistemology: Essays in Honour of Plantinga's Theory of Knowledge*, J.L. Kvanvig, ed. Lanham, MD: Rowman & Littlefield.
- \_\_\_\_\_. 1997. 'How to Resolve the Pyrrhonian Problematic.' *Philosophical Studies* **85**: 229-49.
- \_\_\_\_\_. 2004. 'Replies.' In *Ernest Sosa and His Critics*. J. Greco, ed. Malden, MA: Blackwell.
- \_\_\_\_\_. 2009a. 'Replies to Brown, Pritchard, and Conee.' *Philosophical Studies* **143**: 427-40.
- \_\_\_\_\_. 2009b. 'Replies to Commentators on *A Virtue Epistemology*.' *Philosophical Studies* **144**: 137-47.
- Zagzebski, L. 1996. *Virtues of the Mind*. Cambridge, UK: Cambridge University Press.

