

Frictional Coherentism?

A Comment on Chapter 10 of Ernest Sosa's *Reflective Knowledge*

CRISPIN WRIGHT

Chapter 10 of Sosa's important new book provides an exemplary presentation and discussion of a great dilemma for epistemologists—I'll call it simply the Dilemma. Here is Sosa's statement of it:

Consider the following proposition

KR: a potential knowledge source K can yield knowledge for S only if S knows that K is reliable.

If we *affirm* KR, we face the problem of vicious circularity. How can we attain the required knowledge that our epistemic sources are reliable? Must we not *have* that knowledge *already* before the sources can deliver it to us? How can we know perception to be reliable, for example, without basing our belief on our empirical knowledge derived ultimately from perception? And the same goes for memory and other sources.

If we *deny* KR, that apparently enables us to bootstrap our way from the deliverances of a source on some occasion (or upon a series of them) to conclusions about the safety of its operation on that occasion (or about its general reliability). But this is clearly unacceptable.¹

Since Stewart Cohen's landmark formulation,² the phrase, "The Problem of Easy Knowledge", has most often been associated with the reasoning of the second horn, and has been conceived as a difficulty specifically afflicting a spectrum of externalist and internalist views about epistemic warrant, from various forms of reliabilism to James Pryor's perceptual Dogmatism, which all allow that the acquisition of knowledge by a thinker, S, need involve no policing by S of certain conditions — par excellence, the reliability of the methods or faculties involved — whose failing to obtain would undermine that acquisition. According to such views, I can get to know that P — or more generically, as I shall usually say, get *warrant* for P — just courtesy of the operation of faculties of mine which are reliable in the circumstances in which they are operating, or just courtesy of the (veridical) appearances. I don't have to *monitor* the reliability of my sources, or the suitability of the circumstances for their effective operation, or the obtaining of conditions whose failure to obtain would defeat the appearances. But in that case, on such views, it suddenly appears straightforward to verify at least some relevant such conditions after the event, as it were, without additional epistemic work. Let $P_1 \dots P_n$ be a range of propositions for which I have indeed acquired warrant in the unmediated manner allowed by such views. Since I will also, presumably have warrant for the corresponding n propositions of the form:

¹ P. 211. All page references are to Chapter 10 of *Reflective Knowledge* unless otherwise stated.

² Stewart Cohen, "Basic Knowledge and the Problem of easy Knowledge", *Philosophy and Phenomenological Research* 65, 2 (2002), pp. 309-29.

It (perceptually) appeared to me that P_k ,
 or: The source attested that P_k ,

and since warrant is adjunctive over conjunction, I will thereby be warranted in accepting the n conjunctions of the form

It appeared to me (perceptually) that P_k , and P_k ,
 Or

The source attested to me that P_k , and P_k ,
 and hence in accepting the conjunction of them all in turn. And now, what better evidence could I have for the reliability of the relevant source, or of the appearances delivered to me, than that?

But Sosa is quite right to stress that the problem is not merely an awkwardness for certain kinds of externalist or liberal view but a comprehensive dilemma. For if we react to the foregoing by insisting, with the conservatives, that warrant acquisition must, after all, involve monitoring by the warrantee of (A range of? How many?) enabling conditions and potential defeaters, then it appears we must run into a whole range of problems of which the specific issue – the threat of circularity – highlighted in Sosa’s own formulation is but one. It is a very good question how, in the round, we might accomplish warrant for the judgement that our perceptual faculties are broadly reliable without ultimate reliance of those very faculties. But the concern, more generally, is with the apparent implication of affirming KR, or any monitoring requirement in the same general spirit, that knowledge of anything requires knowledge of the satisfaction of enabling conditions of that knowledge— so that while those who cannot police their knowledge/warrant cannot get any, those who can have to accomplish some kind of *supertask* before they can get any! In general, it isn’t a great idea to make it a condition on the acquisition of warrant that one first acquires another warrant.

Sosa’s chapter grapples with this problem. I’ll turn to his positive proposals about it (in fairness, only broadly sketched in the final part of the chapter) in the third and concluding part of my remarks. But to begin with, I want to outline a response to the Dilemma that seems not to feature at all in Sosa’s landscape; then, in the second part of my remarks, I want to fuss a little about the reasoning of the second —‘easy knowledge’ — horn.

I

First, let me introduce a very broad notion of a *cognitive project*. A cognitive project is defined by a pair: a question, and something one might competently do in order to answer it. Thus there is a cognitive project associated with the question “What time is it”, which one can execute by looking at one’s watch; another cognitive project associated with the same question which one can execute by asking a stranger in the street. There is a cognitive project associated with the question, “What did you have for breakfast”, which one can execute by an operation of memory. There is a cognitive project associated with the question, “Is there a black hole in that region of the sky”, which one can execute by the operation of a radio telescope and a suitable interpretation of one’s findings. Next, let a *presupposition* of a given cognitive project be any proposition expressing a condition doubt about which would rationally

require doubt about the efficacy of a proposed method of executing the project, or the significance of its result. That the clock has not stopped, or that my informant is trustworthy, are thus presuppositions of the two mentioned projects of telling the time; that my memory is functioning reliably is a presupposition of the project focused on my breakfast menu; and that my radio telescope is functioning properly, as well as a whole load of theory about electromagnetic radiation, are presuppositions of the project about the black hole.³

The reliability of a putative knowledge source will, naturally, be one, conspicuous presupposition of any project involving the acquisition of knowledge or warrant via that source. But now, equipped with this notion of presupposition, we can generalise the Dilemma. Can one obtain warrant for a proposition in a way that involves presuppositions $P_1 \dots P_n$ if one does not have independent warrant to suppose that each of these presuppositions is true? If we say “No”, then since the acquisition of warrant by any method whatever will always involve various presuppositions, we are on the brink of committing ourselves to the absurd thesis that the acquisition of warrant for a particular answer to a question always rests on the acquisition of warrant for an answer to a different question; infinite regress or eventual vicious circularity then beckons. But if we say “Yes”, then easy knowledge is threatened by the reflection that, without too much difficulty, we can find logical consequences of almost any proposition, P , that articulate presuppositions of certain specific ways of acquiring warrant for P . Dretske’s original putative counterexamples to Closure of knowledge/warrant across (known) entailment⁴, for example, are of exactly this kind: that the wall in question is not cleverly but deceptively illuminated with coloured light in such a way as to appear to have a different colour to its actual colour is, in the relevant sense, a presupposition of the project of determining its colour by looking; and that the zoo keepers have not been engaged in the fraudulent practice of artfully disguising mundane animals so as to have them, at least to the inexpert, present as indistinguishable from animals of a more exotic kind, is a presupposition of the project of identifying animals in the zoo by their appearance. So if we grant that it is possible to get warrant, or knowledge that a wall is red, or that certain animals are zebras, just by looking, then the view that one needs no warrant for the respective presuppositions in order to acquire these warrants, seems committed — absurdly — to the idea that one can get warrant for the relevant presuppositions just by getting it, purely visually, for the parent propositions and reasoning to the presuppositions as very immediate conclusions. That’s too easy. You cannot tell by looking that there are no tricky conditions in place that preserve the normal visual appearances; and if you cannot tell it by looking, you cannot tell it by looking and kindergarten inference either.

Dretske’s own response to the problem, forty years ago, was to reject closure of epistemic warrant over known entailment. Sosa does not consider this response in his present discussion. Maybe, like many, he considers it desperate.⁵ But there is another

³ Sosa himself makes use of a notion of presupposition throughout the chapter, but I do not know whether the gloss just placed on it would be acceptable to him and will make no assumption about that.

⁴ Fred Dretske, “Epistemic Operators,” *Journal of Philosophy*, 67 (1970), pp. 1007–1023.

⁵ His defense of safety-based conceptions of knowledge in “Relevant Alternatives, Including Contextualism”, *Philosophical Studies* 119 (2004), pp. 35–65, leans quite heavily on the claim that safety corrals most if not all of the intuitions about cases explained by sensitivity while avoiding

response that he does not consider, though it comes very close to the surface in the following passage:

... might the following two things be true conjointly?

- (a) KR is false: in particular a potential knowledge source K *can* give one knowledge that the wall one sees is red, despite ones *not* then knowing that K is truth-reliable
- (b) One cannot know the wall to be red by taking ones visual experience at face value unless (i) in believing the wall to be red one *presupposes correctly and justifiedly* that ones situation is trustworthy for judging whether the wall is red through taking the look of things at face value, and (ii) *in so presupposing one is not doing so arbitrarily, or in some other epistemically defective way.*

The crucial question now concerns the relation between *this* epistemic accomplishment, the one in (b), and knowledge that ones source K is reliable when, on the basis of its look, one believes the wall to be red. Is such ‘presupposing correctly and justifiedly’ a form of *knowing*?⁶

The intuitive idea being canvassed here is that one might steer one’s way between the horns of the easy knowledge dilemma by making something of the idea that although KR is false, – one does indeed not need to *know* of the reliability of a source before it can give one knowledge, – we don’t open the way to easy knowledge by this admission because the source will indeed bestow knowledge on an agent only in a context where the agent already “correctly and justifiedly” *presupposes* its reliability, and that by being thus presupposed, the reliability of the source – or indeed the failure of any relevant defeater – is banished from the range of propositions for which the source can bestow warrant. That there are no operative tricks of lighting is a presupposition of the warrant acquired visually for the claim e.g. that the wall is red; so that visual warrant accordingly does not extend to it. But the requirement is that it be presupposed, not that it be antecedently known.

Reasonably enough, Sosa proceeds to scrutinize the requirement that the presupposition be “correct and justified”, that it is not made “arbitrarily or in some other epistemically defective way”. He allows that such presupposing can be intelligibly taken to be “implicit”. But he quickly rejects the proposal as doomed in any case, on the ground that legitimate presupposition will still have to be a kind of “quasi-knowledge”, so that the propositional attitude that constitutes it is

... one that requires its own positive normative standing in order to do its proper epistemic work, a standing subject to the traditional epistemic framework involving: rational basis, Gettierization, truth, defeat, challenge, defense, competence, aptness, etc.⁷

generating problems for Closure. Both claims are of course challengeable. For my own part, I am inclined now to think that the issues about the validity of Closure are much more fraught than I once supposed, but much of the contemporary discussion unfortunately clouds them by treating Closure as a principle about feasible knowledge *acquisition*— thus, in effect, confusing it with Transmission.

⁶ Page 221.

⁷ Page 222.

In short, presupposition doesn't help since it has to be *justified* presupposition, and this requirement then takes us into a space of assessment of presuppositions which has no advantage over the flat-out endorsement of generalized KR, that legitimate presuppositions have to be known.

But Sosa makes an assumption here. It is one thing to grant that a presupposition – of e.g. the reliability of appearances – has to be justified. It's another thing to grant that this requirement of justification takes us into the space of “rational basis, Gettierization, truth, ...” — in short, into the space of *evidence and cognitive achievement*. To move straight from the first to the second is to pass over the possibility that the warrant required of a presupposition, if it is to mediate the acquisition of knowledge from a source, may allowably be of a different — *non-evidential* — character.

The idea of a non-evidential form of warrant is associated with the terminology of epistemic “entitlement” as it has featured in recent work of Burge, Peacocke, others and myself.⁸ This is not the time or place to try to elaborate and defend such a notion. But in rough outline, and on my own conception of the matter,⁹ the presuppositions to which one is entitled all articulate conditions which, in the course of a particular cognitive project, one may rationally trust in, or take for granted, precisely *without any requirement of cognitive work*. Trust is, in the nature of the case, something one does without evidence; if you have evidence that a source is trustworthy, you don't need to trust. There are a variety of ways of arguing that trust may nevertheless be a rational stance, or frame of mind. Those that I myself find most impressive flow from the observation¹⁰ that *all* enquiry involves so far untested presuppositions, some specific to the particular enquiry at hand, others generic and recurrent; and that the attempt to improve one's epistemic position in this respect is doomed to failure, either because counterparts of the original presuppositions recur or because they themselves recur as presuppositions of the second investigation (as for example when we find ourselves reasoning in ways which presuppose Modus Ponens in attempting to justify that very rule.)

All this is of course controversial and needs elaboration. But assuming its general direction is correct, the idea that an *ideal* system of belief can be coherently conceived

⁸ See, for instance, Tyler Burge, “Content Preservation”, *The Philosophical Review* 102,4 (1993) pp. 457-488, and “Perceptual Entitlement”, *Philosophy and Phenomenological Research* 67, 3 (2003), pp. 503-548; Christopher Peacocke, “Explaining Perceptual Entitlement”, Chapter 3 of *The Realm of Reason*, Oxford: Oxford University Press, 2004; and Wright, “Warrant for Nothing (and Foundations for Free)?”, *Aristotelian Society Supplementary Volume* 78 (2004), pp. 167–212.

⁹ Care is needed here over ‘non-evidential’. Both Burge and Peacocke are writing within preferred epistemological frameworks which, for their own reasons, war with the description of perception, for instance, as a source of *evidence*, strictly and properly so regarded. Evidence, for these authors, is essentially a matter of what is independently known or justifiably believed. That is consistent with regarding one's perceptual state as conferring knowledge that P, or at least as enhancing the likelihood that P. So perceptual entitlement, for Burge and Peacocke, is still an evidential form of warrant in a looser, more general sense of ‘evidence’: perceptual knowledge is achieved by dint of one's entitlement to accept the ‘evidence of one's senses’, and acceptances to which one is so entitled will, when true, count as knowledgeable on that account. My own conception is crucially different, as the discussion following in the text will make clear. ‘Entitlement’, for me, though it mandates *rational* acceptance, has no direct connection with knowledge, or likelihood of truth

¹⁰ Recurrent in Wittgenstein's notes *On Certainty*

as one where each component owes its place to a specific cognitive feat in which, for the first time, we move from ignorance about it to a recognition of its truth, or probable truth, is deeply misconceived. Exactly that, of course, is the ideal characteristic of *foundationalism*. To reject it is to acknowledge that all cognitive achievement – at the least, all *reflective* cognitive achievement in the sense important to Sosa – takes place in a context of acceptances that are mandatory if the achievement is to be rationally claimed but which are not themselves the product of any specific cognitive achievement to date. This is indeed a kind of epistemic predicament, and to one in the grip of the foundationalist ideal, it will impress as tragic and a motivation for scepticism. But the right response is rather that the foundationalist ideal radically misunderstands the nature of epistemic warrant. That response may seem to take a step towards some kind of coherentist conception of justification — something of which Sosa shows himself supportive. However I think that the question whether coherence — whatever exactly it may be taken to consist in — is of itself a *justification*-conducive factor, is another, further issue.

If there is merit in these ideas, then we can glimpse the prospect of a satisfying resolution of the Dilemma along the following lines. Reflective knowledge of a proposition achieved by execution of a particular relevant cognitive project does indeed require warranted acceptance of the presuppositions of that project. So there is no question of coming to easy knowledge by inference that such presuppositions are met; rather they have to be in place already to mediate the acquisition of warrant for the intermediate propositions (that the wall is red, or that the animals are zebras) that supply the premises for the relevant inferences. In short, the easy warrant inferences involve a failure of *transmission* of warrant, as I have elsewhere termed it. But since these presuppositions can be rationally trusted without additional epistemic work, no circle, or regress, of further justification is thereby launched. So we do deny KR: It is not a necessary condition for the derivation of knowledge from a particular knowledge source that one *know* that the source is reliable. That denial does not, however, set up the second — easy knowledge — horn of the dilemma, since we affirm a qualified form of KR: A potential knowledge source K can yield knowledge for S only if S either knows *or may rationally trust* that K is reliable.

II

Sosa speaks in terms our “bootstrapping” our way from the deliverance of a source on some occasion—or from *a series of them*—to conclusions about the safety of its operation on that occasion—or *about its general reliability*. There are reasons for doubting that so clearly illicit a general conclusion really threatens any KR-rejecting view.

This was the envisaged argument. Suppose I receive a series of pieces of information from what is the matter of fact reliable source, all of which are true and qualify as knowledge/warranted by the canons of a view, perhaps Reliabilist, that is supposed to be in difficulty on the second horn of the Dilemma. I am then in a position to know an indefinite number of conjunctions of the form:

The source delivered to me that P, and P.

If I know of no case where the source has let me down, I will thus apparently have a substantial *track record* of successful performance by the source in question, and no evidence for its unreliability. And from there the inference seems invited that the source is indeed generally reliable, notwithstanding the strong intuitive impression that I am in no position to conclude any such thing. But the inference may seem to be objectionable for a reason which even views which propose that KR is false, and that the judging subject need stand in no particular epistemic attitude to the proposition that the source is reliable, are able to put forward.

Some care is needed with the formulation of the relevant point. Here is how I put it in another discussion:

A pool of evidence should be regarded as providing inductive confirmation of a hypothesis only if its reasonable to consider it is drawing upon a representative sample. And that in turn requires a significant prior probability for the thesis that counterexamples would have shown up in the sample if there were any. But the body of “confirming” data compiled by chalking up pairs of the form, <The source has delivered it to me that P; P> in the way described has no chance of containing any counterexamples to the contention of the reliability of [the source]. So it provides no inductive support for that contention in any case, for purely general methodological reasons.¹¹

However, as Stewart Cohen pointed out to me, this is incorrect. It may well be that I do have good reason to think that the sample of cases over which I ‘test’ the source is representative, and that the likelihood that it would contain counterexamples to the source’s reliability, if such exist, is high. What *is* true, however, is that, so long as my sole evidence concerning the truth of a verdict issued by the source is the source’s own word for it, so to speak, I won’t be in a position to *detect* the occurrence of any counterexamples. Since I know in advance that I won’t be a position to detect any counterexamples should they occur, — that I will falsely believe of any counterexample that does occur that it is not one — it seems plainly irrational to regard my “findings” as a confirmation of reliability.

This train of thought appeals to something like the following principle:

One is justified in taking a body of evidence as confirmatory of a given hypothesis only if, were it to be flawed in ways that would prevent it from being confirmatory, one would be in a position to detect the fact.

While I am doubtful that *any* principle of such simplicity and generality holds good in this territory without qualification, it does seem intuitively to be no worse than an oversimplification of something correct. And if so, it will be open to Dogmatism, for instance, to confront the original generalised form of the easy knowledge argument head on and dismiss it by appeal to such a (refined) principle without any compromise of its characteristic denial of KR.

¹¹ At pp. 43-4 of Wright, “The Perils of Dogmatism” in *Themes from G. E. Moore: New Essays in Epistemology and Ethics*, edited by S. Nuccetelli and G. Seay, Oxford: Oxford University Press 2007, pp. 25-48.

Interestingly, though, it is less clear that pure externalist views will have the same recourse. The response is implicitly trading on a distinction, in the way it is handling the notion of ‘detection’, between the *actual character* of my evidence — which, in the case where the source really is reliable, is indeed knowledge of each of the conjunctions, <The source delivered to me that P_k , and P_k >, — and its *subjective character*, whereby it is allowed that for all that is apparent ‘from the inside’, as it were, I might fail to have that very body of evidence; that there are possible circumstances, subjectively indistinguishable from those I find myself in, in which I would not have it. But more than that: the response is attaching *epistemic significance* to this distinction — it is asserting that a hypothesis should not be regarded as having been properly tested unless its failure under test would be in the relevant sense detectable by, that is, would show in the subjective character of the experience of the tester, in what she would be aware of in the course of the test. If that is right, then this way of responding to the generalised form of the easy knowledge argument is not available to straight reliabilist, or other forms of pure externalist conceptions of knowledge and warrant, whose cardinal thesis is in effect the justificational irrelevance of the subjective. For these views, the challenge to address the generalised form of the second horn of the Dilemma may remain.¹²

But be that as it may, the issue is in any case a distraction in the present context. For there is surely already an easy knowledge problem *before* the collation of a ‘track record’ of reliability and its notional generalisation: a simpler form of bootstrapping reasoning which will serve quite well enough for the second horn of the original Dilemma. When a reliable source delivers the true proposition that P , then — also knowing, as I do, that the source has delivered that P — it seems, on the KR-denying views in question, that I may infer the conjunction as before and hence, whatever I may or may not be in a position to know about the source’s *general* reliability, thereby apparently come to know at least that it proved accurate *on this occasion*. And that conclusion, though more modest than that of general reliability, still seems to overreach in much the same way as the conclusions about the absence of cleverly disguised mules and of deceptive red lighting. Intuitively, it is only in a context where we are *independently* warranted in taking it that the source will be accurate, at least in the particular case, that we can take ourselves to be warranted in the transition from the deliverance of the source, whatever it is, to the conclusion that it is true. But this conservative thought returns us to the first horn of the Dilemma.

It thus seems fair to assume that, whatever the outcome of the skirmish with respect to the inference of general reliability, the second horn of the original Dilemma retains its point. To put it at its most general: the thought that the acquisition of epistemic warrant can in some systematic way dispense with the policing of presuppositions of the reliability of the cognitive processes involved appears to come with the cost that the satisfaction of certain particular such presuppositions can be verified just by

¹² I have heard it suggested that an independent objection to the bootstrapping inference to general reliability is furnished by the aggregation of risk of error involved in compiling the conjunction that expresses the source’s ‘track record’. But the observation that the risks of error aggregate, though correct, is not to the purpose. The same is true in ordinary cases where a genuine track record of reliability is compiled, involving independent checks on a source’s performance. That is something which sound inductive methodology has to reckon with in any case. It provides no basis for an independent criticism of the bootstrapping move.

inferring it from beliefs generated by reliance on those very cognitive processes, without any independent check. And that continues to seem preposterous.

III

I have noted two ways of addressing the Dilemma: a denial of KR, adjoined to a rejection of closure of warrant across (known) entailment; and the warrant transmission failure/epistemic entitlement package canvassed in section I above, which likewise denies KR but accepts that reliability of sources does at least need to be rationally presupposed. Sosa's own response is neither of these, and is briefly suggested only in the final pages of the chapter. My grasp of it, and attendant misgivings, are speculative and may well reflect misunderstandings.

Sosa writes:

The right model for understanding reflective justification is *not* the linear model whereby justification is a sort of liquid that flows through some pipe or channel of reasoning, from premises to conclusion. (Such flow is linear, unidirectional; the pipe or channel “transmits” the justification – or warrant, or epistemic status.) A better model is rather that of the web of belief, whereby the web is properly attached to the environment, whilst its nodes can also gain status through mutual support. Any given node is thus in place through its connections with other nodes, but *each of them* is itself in place through *its* connections with the other nodes, including the original given node.¹³

On first reading, this may come across as an elegant expression of the kind of coherence conception of epistemic warrant familiar from Davidson's work but earlier advanced in the writings of the British Idealists and ultimately deriving from Hegel. But in Sosa it cannot be exactly that. We need to make a distinction here between two modes of rational believing: *access rationality* and *management rationality*. A belief is *access-rational* when it is entered into for sufficient supportive reason – when the In-rules, as it were, for beliefs of that type, together with other aspects of the subject's epistemic state, mandate his coming to that particular belief. Dogmatism, for example, is thus a thesis about what it takes for perceptual beliefs to be access-rational. A belief is *management-rational*, on the other hand, when it is properly situated in an organised system of belief, exhibiting certain holistic virtues: consistency, systematicity, explanatory equilibrium, predictive power and simplicity, are included in the usual lists. Coherentism— I should like to say, “Coherentism Proper” or “Pure Coherentism”— should, I suggest, be understood as advancing the following distinctive thesis: that rational access is always *intrasystematic* access: in effect, indeed, that there is no such thing as access rationality for a belief which is formed other than by inference, whether deductive, inductive or abductive, within an *up-and-running* system of beliefs. The raw materials for a dynamic such system — the continuing influx of immediately formed perceptual beliefs, e.g., — are rational only at the point when they have successfully been integrated into the system in ways that preserve its systematic virtues.

¹³ Pages 239-40

Notably, there *is* no easy knowledge problem for pure coherentism of this stripe. The Dilemma is precisely a problem about how to understand the justificational architecture of access rational beliefs independently of considerations of their systematic integration into a larger system. The ground is cut from under the Dilemma by the pure coherentist's willingness to disavow that there is any such well-conceived species of access rationality. So Sosa's view, although his recourse to the metaphor of the web emphatically commits him to coherence as *a* source of justification, cannot be such a pure coherentism.¹⁴ He could hardly have taken the Dilemma so seriously if it were. And the game is given away in any case by his allowance that the web should be "properly attached to the environment"; that allowance requires that there be beliefs whose good standing derives from some other root than their integration within a coherent system.

The key to understanding this notion of "proper attachment" is presumably Sosa's conception of animal knowledge. Animal knowledge does not require any policing. It is pure reception of information, exactly a kind of knowledge that a subject can have without any knowledge of the reliability of the sources that give it to him. Reflective knowledge, on the other hand, will be distinguished, presumably, by requiring knowledge of the reliability of its sources and indeed of its presuppositions generally. How is that to be accomplished, in Sosa's view?

Sosa's tendency seems to be to try to address the Dilemma by, as it were, distributing the distinction in the two kinds of knowledge across the horns. So: there has to be some knowledge (the animal) for which KR fails if knowledge is to be possible at all; but there also has to be some knowledge for which KR holds if knowledge is to allow rational scrutiny and organisation under a fully responsible epistemic perspective. Yet it is difficult to see how the distribution helps. For we can simply present the Dilemma as one for reflective knowledge. How are we to advance from the animal knowledge that the wall is red to reflective knowledge that it is so? If reflective knowledge requires reflective knowledge of the presuppositions of its acquisition, how is the latter to be accomplished?

An *impure* coherentist answer is possible—in effect, a pure coherentism restricted to *reflective* knowledge and warrant. Call it *frictional coherentism*. On the frictional coherentist picture, we find ourselves smitten, as it were, with a continuing influx of (perceptual) basic knowledge—animal knowledge—and we elaborate a coherentist web around these animal-knowledgeable beliefs in which the presuppositions of taking them to be knowledgeable are held in place purely by their coherence within the system as a whole. The rationality of the presuppositions that support reflective knowledge of the outcomes of our perceptual interactions with the world flows, accordingly, not in general—impossibly—from prior specific investigations of those presuppositions nor from non-evidential entitlement, but from the coherence of the overall system of belief in which they are anchored: the system which issues licenses for them in specific instances.

That is the kind of thing a reader might well expect Sosa to say. But it isn't evidently what Sosa wants to say—in fact, he nowhere explicitly entertains such a proposal.

¹⁴ And is much the better for that, many would hold. It is pure coherentism that is open to the McDowellian complaint about "frictionless spinning in the void".

The clearest statement I can find concerning how it seems he would like to think about the matter is as follows

Consider one's justification for a given commitment (or its status as epistemically appropriate): say a commitment that lies behind one's belief that one sees a red wall

— for instance, the commitment that in these circumstances the appearance of the wall displays its actual colour —

- might one's *reflective* rational justification for that commitment gain a boost through one's now basing it in part (perhaps in some very small part) on the belief that one does see a red wall? How are we to understand such boost in reflective rational justification?¹⁵

I take the positioning of this passage — right before the final, web-versus-pipeline remarks — to indicate that Sosa conceives of the web model as doing something different from, or additional to, providing for a kind of top-down, coherence-based reflective justification for the presuppositions of a cognitive project. The web model is somehow to provide for a “boosting” in one's reflective rational justification for a presupposition by the very belief whose formation, in the relevant circumstances, it underwrites. So the idea seems to be that once a suitable web of beliefs is in place, able to receive and integrate new animal knowledge and take it up into reflective awareness, a *modest degree* of transmission of warrant – not enough, presumably, to invite the charge of “easy warrant” – from immediately believed (animally known) premises to those of their deductive consequences that articulate presuppositions for the acquisition of that knowledge, becomes possible.

How would that work? Let a normally sighted agent A stand before a red wall, and let it look red to her. And let it be that its appearance does in the circumstances provide her with unreflective—animal—knowledge that (*Red*) the wall is red, notwithstanding her lack of any reflective warrant to think (*P*) that that in these circumstances the appearance of the wall displays its actual colour. Suppose A's unreflective knowledge of (*Red*) surfaces in a self-conscious belief to that effect. If, now at the level of reflection, A is to credit herself with a warrant, acquired just by looking, for this belief, she will need to presuppose (*P*). And this presupposition will need to be in good reflective rational standing if her belief in (*Red*) is to be so. Does Sosa's suggestion about a “boost” point to an explanation how the presupposition of (*P*) can *get to be* in good reflective rational standing, and thereby underwrite reflective knowledge of (*Red*)?

I cannot see that it does. To begin with, the principle seems compelling that a rational agent cannot apportion *more* confidence in the outcome of a cognitive project than she has in anything she recognises as a presupposition of it. So A cannot rationally apportion more confidence in (*Red*) than the degree of confidence she has in (*P*). Hence if she starts from a portion of rational open-mindedness about (*P*) and undergoes only a “very small” positive shift from that in coming to the belief that she sees a red wall, only a very small degree of resulting reflective confidence in (*Red*) is going to be rational— not enough, presumably, to attain to reflective knowledge of (*Red*). Second, there is a concern that Sosa's suggestion tacitly switches the epistemic

¹⁵ Page 239.

operator. Even if the transmission of a modest degree of reflective rational confidence across the entailment from her beliefs about the character of her present experience and the colour of the wall to (*P*) is granted, these beliefs must already enjoy some measure of *reflective rational* confidence if there is to be any, even a “very small” quantity, of that to transmit to (*P*) in the first place. And finally, even the suggestion of transmission of a very small degree of rational warrant offends the original intuitions about the case. A claims to be rationally confident about the colour of the wall on the basis of how it looks to her. We ask her, “What reason do you have for thinking that the circumstances are such that the appearance of the wall displays its actual colour?” She replies, “Well, nothing very strong; but there is the appearance of the wall.”

I am, though, far from (reflectively rationally) confident that I have understood Sosa as he intended. So let me not pursue these concerns but close instead by tabling three questions:

- Does a frictional coherentist response to the Dilemma—consisting in an endorsement of the distinction between animal and reflective knowledge, a denial of KR for the former, a coherentist account of reflective knowledge, and an acceptance of KR for the latter—resonate with his own thinking, at least as a general direction?
- If so, what is the role of the “boost” suggestion within such an account?
- If not, might he further explain that suggestion and how it might be worked into a response to the Dilemma?

New York University
The Northern Institute of Philosophy, University of Aberdeen