

EVERYBODY NEEDS TO KNOW?

This reader came away from Sosa's *Judgment and Agency* with the poignant impression of an otherwise sophisticated and compelling view encumbered by an implausible central element. The sophisticated and compelling view is Sosa's 'virtue reliabilism'. The implausible element is the appeal to the notion of an 'intention to affirm aptly'. These comments explore the possibility of an unencumbered virtue reliabilism in which the role Sosa assigns to such intentions is played by something else instead.

The discussion has four parts. §1 sketches just enough of the *Judgment and Agency* system to bring out the role intentions to affirm aptly play within it. §2 introduces the notion of what I shall call a 'need' – a mental state which is like an intention in that it moves us to action, but unlike an intention, is not a propositional attitude. §3 explores the possibility of a form of virtue reliabilism in which appeal to such a need – the need to attain knowledge – replaces Sosa's appeal to intentions to affirm aptly. §4 raises some questions that might be explored in comparing Sosa's framework (intention-based virtue reliabilism) and the need-based alternative proposed in §3.

1 Intention-based virtue reliabilism

This section sketches Sosa's virtue reliabilism. It is not possible in so constricted a space to explain the view in its full sophistication. So I shall say just enough to isolate the element upon which I wish to focus: the appeal to intentions to affirm aptly.

Here is what I take to be Sosa's central suggestion. Forming a belief or making a judgment is a performance aimed at an outcome. And an account of the normative status of this kind of performance should proceed along the lines we use in assessing normative status for performances-aimed-at-outcomes in general. In general, when assessing whether such a performance has gone well, we ask the following three questions:

- 1 What was the outcome at which the performance aimed?
- 2 Was this outcome attained?
- 3 Was it attained through the exercise of the subject's competence in executing such performances (in a way that *manifests*¹ this competence), rather than coming about as a matter of luck?

If the answers to 2 and 3 are 'yes', the performance reaches the gold standard as far as positive normative status for performances is concerned: it is a performance whose aim is achieved through the exercise of the subject's competence. Propositional knowledge is what we get when a performance of belief formation reaches this gold standard. Belief formation is a performance that aims at truth. A belief counts as knowledge iff it is true,

¹ The notion of 'manifestation', carefully explained in Ch2, is a primitive in the framework. The suggestion is that a performance manifests a competence iff it is causally derived from the competence in a way that involves no deviant causal chains, where the right hand side of this biconditional is not to be regarded as explanatorily prior to the left: causal derivation of performance from competence without a deviant causal chain is just what there is in cases of manifestation.

and the subject's arriving at a true belief in the instance in question is a result generated in an appropriate way by the subject's exercise of belief-forming competence.²

When combined with a suitably generous account of what it is to form and hold a belief, this proposal generates a quite catholic account of what can count as propositional knowledge. Suppose that there is belief wherever there is 'aiming' at truth, and there is competence in belief formation wherever beliefs are formed by a means that is reliable at securing this aim. Then a subject can count as a believer just in virtue of having states formed by means that in some sense 'aim' at truth. So my dog has beliefs. And a state of a subject can count as a state of knowing in virtue of being an 'aims at truth' state; being true; and having its truth secured in an appropriate way by the reliable means of formation deployed in arriving at it. So my dog knows things.

If we just stopped there, the normative status of those among my beliefs which count as knowledge would come out as no different from that of various states of my dog. This would be at least a disconcerting result. But Sosa does not just stop there. He distinguishes, within the over-arching framework, two kinds of positive normative status for performances:

Animal aptness – a performance is 'animally apt' iff it is aimed at bringing about some end state that can be specified without mention of how it is to be brought about, and this aim is attained through the exercise of the subject's competence in executing such performances.

Full/reflective aptness – a performance is 'fully apt' iff it is aimed at bringing about some end-state as an appropriately generated result of the exercise of the subject's competence in executing such performances, and the end-state is in fact brought about in this way.

For the epistemological case, the suggestion is now that human knowledge – the traditional subject matter of epistemology – is *fully apt belief*. The mental states of ours that count as 'knowledge' in this sense are 'judgments'. They are beliefs formed with the aim of *being true in virtue of being formed by an exercise of true-belief-forming competence* which do in fact have their truth secured in this way.³

So far we have said nothing about intentions. Intentions enter the framework with Sosa's specification of the aims of performances which are sophisticated enough to be up for evaluation as fully apt or not. Sosa is explicit about the fact that the aims characteristic of performances that have (in the best case) only animal aptness need not be encoded in intentions.⁴ But where a performance is fully apt, he suggests that it has an aim which *is* encoded in an intention.⁵

These elements combine to give the following account of judgment and of our (human) knowledge:

² For summaries incorporating all or some of these elements see 1, 19, 65-77

³ For the distinction between 'full' and 'animal' aptness see 67-73. For the distinction applied to belief-formation as a performance, see 74-77.

⁴ See especially 67-69.

⁵ See for example 166.

1 A subject making a judgment intends to form a true belief as a result of the exercise of true-belief-forming competence.

2 The resulting belief counts as knowledge iff (i) it is true, and (ii) the subject's formation of a true belief in this instance is appropriately secured by (manifests) the exercise of the competence mentioned in 1.

The next three sections explore the possibility of a view which goes along with much of Sosa's discussion, including the claim that human knowledge is fully apt judgment, but disembarks at the claim that the aim of judgment is encoded in an intention.

2 The 'conceptual'/ 'non-conceptual' distinction for the case of motivational states

This section develops the distinction between two kinds of motivational state: intentions, and what I shall call 'needs'.

Let us start by thinking not about motivational states, but about representational ones. I shall suppose – following Peacocke⁶ – that a mental state is 'representational' iff it has correctness conditions. And I shall adopt Peacocke's account of the distinction between conceptual and non-conceptual representational states: a representational state is 'conceptual' iff being in it involves grasp of its correctness conditions, and 'non-conceptual' otherwise. So a mental state is a non-conceptual representational state iff it has correctness conditions, but a subject can be in it without having the capacity to grasp them. Many people have argued that beliefs are conceptual representational states and perceptual experiences non-conceptual ones: you cannot believe that the defenestration of the lords regent precipitated the Thirty Years' War unless you grasp the concepts *defenestration*, *lord regent*, *precipitate*, and *Thirty Years' War*, but a subject could have a perceptual experience that represents a blue cube without grasping the concepts *blue* and *cubical*.

There is obviously room for argument and discussion about the details of this basic framework for the case of representational states. However, let me just ignore the details, and move to the suggestion I want to make: the basic 'conceptual'/ 'non-conceptual' distinction transfers quite readily to the case of motivational states. To see how, note first that, just as a representational state has correctness conditions, a motivational state has what I shall call 'fulfilment' conditions – conditions which must be met if the state is to be fulfilled. For example, if I intend to A, I am in a motivational state whose fulfilment requires my A-ing (or my A-ing in a way appropriately guided by my intention). But having recognised that motivational states have fulfilment conditions, it is a short step to a conceptual/non-conceptual distinction that parallels the distinction we drew in the representational case:

A motivational state is 'conceptual' iff being in it requires grasp of the concepts that would be involved in specifying its fulfilment conditions

A motivational state is 'non-conceptual' iff being in it does not require grasp of the concepts that would be involved in specifying its fulfilment conditions.

⁶ 'A Study of Concepts' Ch3.

With the distinction in place I do not think it is an unwarranted step to the claim that it applies to the motivational states that guide our own actions. An intention is a conceptual motivational state: to have an intention, you have to grasp what it is an intention to do. (If I am to intend to defenestrate the lords regent, I must grasp the concepts *defenestration* and *lord regent*). But we are also moved to action by non-conceptual motivational states – what I shall call ‘needs’. The most familiar needs are animal or emotional. They are needs many of which we share with other higher animals, and which are characteristic of our emotional lives: the need to avoid hunger; the need to avoid loneliness; the need to feel secure. Dogs, small children, and other beings who lack the conceptual sophistication required to grasp the fulfilment conditions of these needs routinely act in ways motivated by them: these are non-conceptual motivational states.

And having recognised animal and emotional needs, there is no apparent barrier to recognition of what I shall call ‘rational’ needs – needs the having of which is characteristic of our rational lives. For example, I think it is plausible that we need (when we are awake) to be sustaining a comfortable level of cognitive activity, and need to have a comfortable level of coherence in our cognitive lives.

So here is a suggestion. One of our rational needs is the need to know. This is a non-conceptual motivational state – a state which guides much of our belief-forming activity, but the having of which does not require the cognitive sophistication involved in specifying what it is a need for.

3 Need-based virtue reliabilism

Now let us consider how the possibility of motivation by a rational need might be grafted into Sosa’s virtue reliabilist framework. The basic suggestion would be to replace Sosa’s appeal to an intention to form true beliefs by means that manifest true-belief-forming competence with an appeal to a need to do the same thing. So instead of Sosa’s 1 and 2 from the end of §1, we would have the following:

- 1 A subject making a judgment is moving to a belief under the guidance of the need to form true beliefs whose formation manifests true-belief-forming competence.
- 2 The resulting belief counts as knowledge iff (i) it is true, and (ii) the subject’s formation of a true belief in this instance in fact does manifest this competence.

Now, this proposal will get off the ground only if we can make sense of the notion of ‘guidance by a need’. And in fact, it is in explaining how the motivational state characteristic of judgment plays its guiding role that it seems to me that Sosa’s claim that this motivational state is an intention does most work.⁷ So one possible point of resistance to need-based virtue reliabilism is right here, at the notion of guidance by a need – an objector might say that a non-propositional motivational state cannot play the kind of ‘guiding’ role that fully apt judgment (knowledge full well) requires.

⁷ The suggestion is that the sophistication of the intention brings with it an extra layer of non-luckiness in the success of the guided performance. See the discussion of pp84-88.

But I think it would be premature to give up on need-based virtue reliabilism in the face of this objection..

One reason not to give up is that if we allow (as surely we should) that some of our agential behaviour is driven by non-conceptual motivational states, we shall require an account of guidance of behaviour by needs *anyway*, regardless of the role this kind of guidance might play in the kind of epistemological project I have sketched.

Another reason is that the rough outlines of an account of guidance by needs is actually ready to hand. The account I have in mind is built around the general notion of guidance by a goal-representing state:

A 'guidance' system is characterised by 1-4:

- 1 The system contains an internal representation of a goal.
- 2 The system also contains a feedback mechanism which generates reports on current status.
- 3 Moves made by the system are determined by the goal represented at (1) and the feedback about status at (2).
- 4 If the system is functioning normally, then, over time, the function from (1) and (2) to output actions described at (3) tends towards reducing the difference between the goal state represented at (1) and the current state reported at (2).

Given this very general characterization, there is a wide variety in 'guidance' systems. For example, a heating system controlled by a thermostat will count as a system of such a kind – in this case, the goal-representing state is the thermostat setting; if the system is functioning normally, the 'output actions' (furnace on, off, up, down) tend to reduce the gap between the temperature at which the thermostat is set and the ambient temperature that the feedback mechanism detects. (The 1-4 model draws on the 'system-property' view of goal directedness from the philosophy of biology.⁸ As far as I know, appeals to this kind of model of guidance were brought into the philosophy of action by Frankfurt.⁹)

Obviously we do not want to say that the thermostat setting is a 'motivational state' of the heating system. So we need to distinguish guidance in general from guidance by motivational states of the subject:

Guidance by a motivational state is the special case of 1-4 where the goal-representing state is a personal level state of a subject.

Given this account of guidance by a motivational state, guidance by intention comes out as one case of the general phenomenon (the case where the goal-representing state is a conceptual state of the subject), and guidance by a need comes out as another (the case where the goal-representing state is non-conceptual).

Of course, even given this general account of guidance by needs, it will remain to show how any particular need does its behaviour-guiding work. For the case of the need to know, the idea will be that the goal and the feedback mechanism steer us towards belief-forming activities which will result in true beliefs that manifest out belief-forming

⁸ See Nagel 'Goal-directed Processes in Biology' §4; Adams 'The Goal-State Theory of Function Attribution' 506-7.

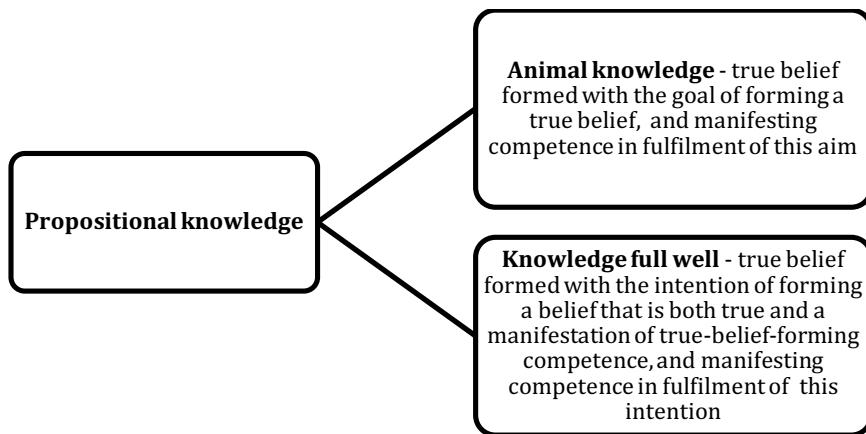
⁹ In 'The Problem of Action'.

competence unless the situation is in some way unusual or devious. (If the situation is devious or unusual, the subject will end up with a case of competent but failed (because unlucky) performance). For example, for the most central cases of ‘knowing full well’ on the basis of perception, a full version of the kind of view I am sketching would explain how the goal represented – formation of true beliefs which manifest true-belief-forming competence – and the relevant mechanisms for detection of current status – perception; the mechanisms that generate the kind of sensitivity to defeaters characteristic of our perception-based belief-forming – combine to generate the familiar ‘default acceptance with sensitivity to defeaters’ pattern. And the full version would tell a parallel story about ‘knowing full well’ on the basis of testimony.

I take it that there would be substantial structural overlap between these stories as occurring in a full version of need-based virtue reliabilism and the stories about knowledge from perception and testimony¹⁰ already there in Sosa’s intention-based view. The main difference would be that the need-based stories would treat the most central cases of sensitivity to defeaters as achievements more primitive than, and not calling upon, the kind of second-order conceptual representation of beliefs, the means to their formation, what it is for a belief to be true, and what it is for a true belief to manifest true-belief-forming competence which, in Sosa’s framework, is always at work when knowledge full well is attained.

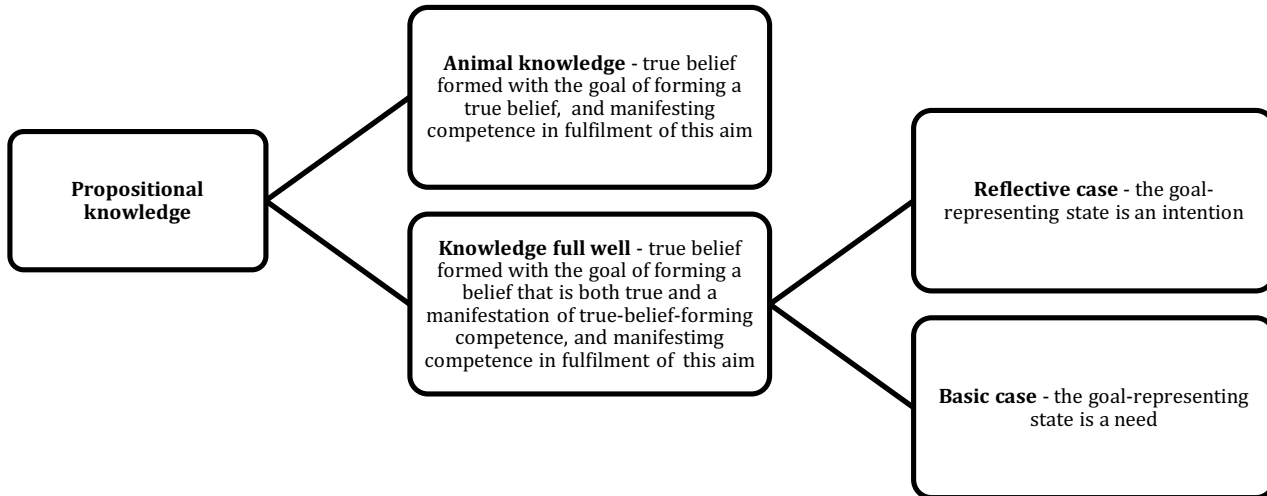
The reader may have noticed the insertion of a hedge in the previous paragraph. I began this section talking about need-based virtue reliabilism as a view of knowledge full well in general. But the claim in the previous paragraph was only that this model should be applied to the most ‘central cases’ of beliefs formed on the basis of testimony or perception. And in fact I do think it is plausible that there is a kind of very reflective knowledge where something like Sosa’s intention-based story applies. Perhaps scientific investigation works like this – the scientist proceeds under the guidance of an intention to form true beliefs that manifest true-belief-forming competence. So the suggestion is in fact that there are three cases where Sosa is distinguishing only two:

Sosa’s view



¹⁰ For the view of knowledge from testimony see for example 116-17. For the view of knowledge from perception see for example 150-153.

The view sketched in these comments



4 Comparison

This final section makes three points of comparison between Sosa’s intention-based virtue reliabilism and the need-based alternative that I have sketched. §4.a considers one central illustration of the puzzle-solving power of Sosa’s framework, and argues that the central features of this part of his discussion are straightforwardly transferable to the need-based framework. §4.b considers a kind of puzzle case for which I suggest that the need-based framework *might* come out ahead of the intention-based one. §4.c steps back to ask which framework looks more attractive considered aside from issues of puzzle-solving power.

4.a Gettier cases and reliable clairvoyance

Sosa spends a substantial portion of *Judgment and Agency* showing how his view generates solutions to traditional puzzles in epistemology. This subsection summarises one strand of such puzzle-solving – the discussion of why our intuitions fall out the way they do with respect to Gettier cases and cases of reliable clairvoyance – and argues that Sosa’s proposal, developed in his intention-based framework, is, in all important respects, available in the need-based framework too.

Consider a ‘fake barns’ Gettier case and a case of reliable clairvoyance:

FAKE BARNS X is in a circumstance whose ordinarieness she has no reason to doubt, and is looking at something which both *looks like* a barn and in fact *is* a barn. X forms, on the basis of her perceptual experience, the belief she would express by saying ‘That is a barn’. However, though the barn-like thing she is looking at as she forms her belief *is* a barn, most of the barn-like things visible from the road along

which X is driving are not barns but barn-facades, erected to encourage the impression that the factory-farmed countryside X is traversing is still cultivated by rustic barn-involving means.¹¹

RELIABLE CLAIRVOYANCE Though Y does not realise it, he is a reliable clairvoyant. Y does not often arrive at beliefs by the exercise of his faculty of clairvoyance. And when he does arrive at beliefs in this way, he is not aware that this is how they have been reached. But, nevertheless, Y does sometimes arrive at beliefs by clairvoyance, and when he does so these beliefs are reliably true.¹²

Starting with these cases, Sosa argues as follows:¹³

i) It is a widely acknowledged intuition that in the fake barn case X's true belief does not count as knowledge. And it is widely held that the reason the belief does not count as knowledge is that it is *somehow* 'unsafe' – there is some respect in which the condition that X could not easily have been wrong in a similar case fails.

ii) It is also a widely acknowledged intuition that the beliefs Y forms by reliable clairvoyance do not count as knowledge because there is *something* wrong with the route to their formation, even though this route is, for Y, a safe generator of true beliefs.

But

iii) Though (i) and (ii) capture one line of intuitive response to the respective cases, there is also an intuitive pull in each case towards the claim that the relevant beliefs *do* count as knowledge.

iv) Virtue reliabilism generates an explanatory clean sweep of the (i)-(iii) data.

So

v) We have an argument from explanatory power for the virtue reliabilist framework.

The case for (iv) hinges on a distinction that the virtue reliabilist framework generates between two notions of 'safety':

Safety with respect to animal aptness ('safety of affirmation'): a belief is 'safe with respect to animal aptness' iff formed by a means which could not easily have generated a false belief.

Safety with respect to full aptness ('safety of judgment'): a belief is 'safe with respect to full aptness' iff formed by a means which is (a) guided by the subject's intention to form a true belief whose formation manifests the relevant true-belief-forming competence, and (b) could not easily have resulted in failure with respect to this intention.

¹¹ Compare Goldman 'Discrimination and Perceptual Knowledge'

¹² Compare Bonjour 'Externalist Theories of Knowledge'

¹³ See 77-82

The second of these safety conditions is much more demanding than the first. Safety with respect to animal aptness requires only that the belief be formed by a means that is reliable across a suitable range of similar circumstances. Safety with respect to full aptness requires that the performance of belief formation fulfil the intention to form a belief in a way that manifests true-belief-forming competence. If the first condition is met but the second is not – if the belief is safe with respect to animal aptness but not full aptness – there is a sense in which it is a mere matter of luck that the route to belief-formation that the subject deploys is a reliable one.

With the distinction between kinds of safety in place, it is a short step to Sosa's account of the (i)-(iii) combination. Sosa argues that both fake barn Gettier cases and reliable clairvoyance involve safety with respect to animal aptness but not full aptness. For the clairvoyant, the second-order intention characteristic of full aptness is missing: Y is not even aware that he has the faculty by which his belief is formed, so can have no intention to form a belief by its means. In the 'fake barns' case the second order intention is present but not fulfilled: X intends to form a true belief in a way that manifests true-belief-forming competence; since the belief is true as a mere matter of luck, this intention is not fulfilled. But in the virtue reliabilist framework a true belief counts as human knowledge/knowledge full well iff it is safe with respect to full aptness. So neither X's belief in the fake barns case nor Y's beliefs formed by reliable clairvoyance count as knowledge full well. However, in each case the beliefs *do* count as animal knowledge: they are true beliefs that are safe with respect to animal aptness. So Sosa can explain both the 'There is not knowledge here' intuitions at (i) and (ii) and the 'There is knowledge here' intuitions at (iii). The 'There is not knowledge here' intuitions are explained by the *absence* of knowledge full well. The 'There is knowledge here' intuitions are explained by the *presence* of animal knowledge.

This strikes me as an elegant account of the target intuitions. But as far as I can see a structural parallel is available in the need-based framework too. The need-based framework can just co-opt Sosa's notion of safety with respect to full aptness, adjusting as required:

Safety with respect to full aptness (alternative version) A belief is 'safe with respect to full aptness' iff formed by a means which is (a) guided by a motivational state (either an intention or a need) whose goal is the formation of true beliefs by means that manifest true-belief-forming competence, and (b) could not easily have failed to fulfill this goal.

We have seen that Sosa's account of the (i)-(iii) combination turns on the claim that the relevant beliefs are safe with respect to animal aptness but not full aptness. But the need-based virtue reliabilist framework generates this same result. The clairvoyant's beliefs are unsafe with respect to full aptness because formed by a means which, though reliable, is not guided by an appropriate motivational state. For the subject in the 'fake barns' case, the mechanism of belief formation (uptake from perception) *is* guided by such a motivational state – the need to form beliefs by means that manifest true-belief-forming competence – but the goal is not met so the need goes unfulfilled.

I conclude that Sosa's virtue reliabilist account of the (i)-(iii) combination transfers quite readily to the need-based framework. The puzzle of how to account for the (i)-(iii) combination is just one of the many for which Sosa has virtue reliabilist

solutions. So it would obviously be a vault from here to the conclusion that the (considerable) puzzle-solving power of Sosa's intention-based virtue reliabilism is shared by the need-based alternative. Therefore let me re-tool this potential conclusion – out of reach in the space available for present purposes – as a challenge: If puzzle-solving power is to provide grounds for preferring intention-based virtue reliabilism to the need-based alternative, there must be puzzles that the intention-based view can solve and the need-based view cannot; the proponent of the intention-based view should tell us which these puzzles are.

4.b Judgment in the face of second-order doubt

§4.a suggested that intention-based and need-based virtue reliabilisms might be on a par with respect to puzzle-solving power. This section raises a kind of case where the two views generate distinct diagnoses, and which might, therefore, provide the starting point for an argument for the conclusion that one of them has puzzle-solving power the other lacks.

Consider the following scenario:

You are having what seem to you to be ordinary perceptual experiences. But a source you take to be reliable has told you that you have been given an hallucinogen, and you have formed a belief to this effect. 'I really shouldn't believe my eyes', you say to yourself. But the experiences that unfold in front of you are so mundane and realistic that you just can't help it. You believe that you are not in a position where beliefs formed by uptake from perception will manifest true-belief-forming competence. But you keep drifting into forming beliefs by this mechanism anyway.

Now consider how Sosa's intention-based virtue reliabilism will treat this kind of case. According to Sosa, judgment requires intending to form a true belief by a means that manifests true-belief-forming competence. But most philosophers who have worked on intention agree that intending to A requires believing that you will A, or is at least inconsistent with believing that you will not or cannot A¹⁴. Given these claims, the belief-formings in cases like the one described cannot be judgments: there can be no cases of judgment in the presence of second-order doubt about whether you are in a situation where judgment is appropriate.

In the need-based framework, this result falls away. Acting under the guidance of a need does not entail believing that you are on the way to fulfilling the need. So the subject in the case described can count as acting under the same motivation as drives ordinary cases of belief formation: as judging in the face of a second order belief that s/he ought not to.

Is this an argument for the need-based model I have sketched? That depends on whether we *should* recognise the possibility of judgment in the face of second order doubts. This is a matter on which I am as yet undecided. I mention this kind of case as one clear instance where the difference between Sosa's intention-based view and the need-based alternative comes to the fore.

4.c A polemic against profoundly implicit propositional attitudes

¹⁴ See for example Anscombe *Intention*; Bratman *Intentions, Plans, and Practical Reason* 4, 15-18; Searle *Intentionality* 408-9; Velleman 'How to Share an Intention' 202-4.

So far in this section I have taken some steps towards comparing the two virtue reliabilist views by their explanatory power. I shall close by saying something polemical in favour of the need-based view.

It is commonplace to allow that only some of a subject's beliefs at a time are 'occurrent' (consciously endorsed by the subject at the time). Others are 'dispositional' (available for conscious endorsement, but not consciously endorsed). It is also commonplace to allow that only some of the things we count as believing are believed 'explicitly' (represented by some specific state of the subject's mind). Other things are believed 'implicitly' (they are not represented by states of the subject's mind, but are swiftly and trivially derivable from things that are). And these distinctions – occurrent vs. dispositional; explicit vs. implicit – seem to apply to other propositional attitudes too. For example, my intention to get these comments in not too long after the deadline is occurrent as I write this sentence, because in writing the sentence I am also stating the intention. But it was merely dispositional a minute ago, and will be merely dispositional a minute hence. And my explicit intention to get the comments in not too long after the deadline entails many implicit intentions of form <I intend not to A tomorrow>, where A is some action whose incompatibility with getting the comments in is swiftly and trivially derivable by me from my explicit beliefs.

In recognizing these distinctions, we allow that a subject considered at a time holds many propositional attitudes that she is not consciously endorsing at the time, and which are not, at the time, even in the stacks available for endorsement – they must first be derived by some swift and trivial step. But there is a tradition of recognising propositional attitudes even more distant from the conscious forefronts of our minds: attitudes we count as holding merely in virtue of the utility of supposing we do in enabling an hypothetical theorist to predict our behaviour or explain its rationality. I shall say that an explanation appeals to 'profoundly implicit' propositional attitudes iff it makes the kind of move I have just described – crediting a subject with propositional attitudes that are neither occurrent; available for occurrence; nor reachable from such attitudes by swift and trivial steps; and whose contents the subject may lack the conceptual resources to formulate.

Though it would be at least difficult, and perhaps impossible to *prove* that appeals to profoundly implicit propositional attitudes are illegitimate, it is easier to show that anyone making such an appeal is keeping bad philosophical company. For the most thoroughgoing constructors of explanations appealing to profoundly implicit propositional attitudes were behaviourists, who denied that we have mental lives at all, and held that that a propositional attitude ascription is only ever a placeholder for an account of regularities between stimulus inputs and behavioural outputs. And the further we move from behaviourism, the less plausible such appeals seem. If you allow that there are facts of the matter about *some* of our mental states, why keep the pseudo-behaviourist claim that all there is to say about some others is what would enable a theorist knowing what the subject perceives and how she behaves to predict or explain her behaviour? In particular, if you abandon the behaviourist view of propositional attitudes which are occurrent, in the stacks, or trivially and immediately derivable by the subject, what non-*ad hoc* grounds can you give for allowing some more subliminal range of propositional attitudes to which the abandoned behaviourist model still applies?

Now, there are various points at which Sosa distances himself from the claim that the intention to affirm aptly must be either occurrent or explicit.¹⁵ But it seems to me that, given the sophistication of this intention, he needs to go further than this. The intention is the intention to form a true belief as a manifestation of true-belief-forming competence. Grasp of this intention's propositional content requires grasp of the concepts of *truth* and *belief*; of *true-belief-forming competence*; and of *manifestation* (with its association with the notion of a deviant causal chain). Is it plausible to think that subjects who lack the conceptual sophistication to grasp these concepts or the proposition constructed from them cannot really make judgments on the basis of perception or testimony? I do not think this is plausible. I do not think Sosa thinks so either. But in that case the intentions at the heart of his virtue reliabilism are profoundly implicit – they are intentions allegedly held and acted upon by subjects who may lack the conceptual sophistication to grasp them. So I suggest that the appeal to intentions to affirm aptly brings Sosa into bad – bad because behaviourist – philosophical company.

References

- Anscombe, G. E. M. *Intention*. Oxford: Blackwell. 1957.
- BonJour, L. 'Externalist Theories of Knowledge'. *Midwest Studies in Philosophy*, 5. 1980.
- Bratman, M. *Intentions, Plans, and Practical Reason*. Stanford, CA: CSLI. 1987.
- Frankfurt, H. 'The Problem of Action'. *American Philosophical Quarterly* 15 (2). 1988.
- Goldman, A. 'Discrimination and Perceptual Knowledge' *The Journal of Philosophy*, 73. 1976.
- Peacocke, C. *A Study of Concepts*. Cambridge, MA: MIT Press. 1992.
- Searle, J. *Intentionality*. Cambridge: Cambridge University Press. 1983.
- Sosa, E. *Judgment and Agency* Oxford: Oxford University Press. 2015.
- Velleman, J. D. 'How to Share an Intention' in Velleman *The Possibility of Practical Reason*. Oxford: Oxford University Press. 2000.

¹⁵ For example at 166-7.