

ARTICLE

Knowing Full Well from Testimony?

Elizabeth Fricker*

*Email: lizzie.fricker@magd.ox.ac.uk

Abstract

Testimony poses a challenge to systematic epistemology. I cite two kinds of testimony situation where the recipient's belief is not safe, yet intuitively counts as knowledge. Can Sosa's more sophisticated virtue reliabilism, which theorises animal knowledge as apt belief, yield the intuitively correct verdict on these cases? Sosa shows that a belief can be apt, though it is not safe, and so it may seem a quick positive answer is forthcoming. However, I explore complications in applying his AAA framework, regarding what we take as the circumstances in which the subject's attempt is made: the AAA framework does not mandate a particular choice, yet this affects whether the attempt (in particular, a believing in the endeavour to attain truth) comes out as apt or not. I conclude that Sosa's theory is subject to a familiar charge: it does not give a reductive account of knowledge, since we must deploy independent intuitions about whether knowledge is gained in a case, in order to apply it.

Keywords: testimony; knowledge; virtue epistemology; knowing full well; virtue reliabilism; Ernest Sosa

1. The problem posed by testimony for a systematic epistemology

One method for epistemology is not systematic: particular sources of knowledge are approached individually, and epistemic principles adduced ad hoc as seems to best explain the data for that source. But many, including Ernest Sosa, are more systematic. Sosa has developed over many years his own distinctive kind of virtue reliabilism. The generality of his theory goes beyond that of subsuming different sources of knowledge under one account: he portrays the normativity of knowledge as one instance of a broader performance normativity – a normativity that accrues intrinsically and essentially to any skilled performance with a constitutive aim. The result is an elegant original theory of considerable explanatory power. My discussion here is based on Sosa's theory as expounded first in his monograph *Knowing Full Well* (KFW) (Sosa 2011), and then refined and extended in his book *Judgement and Agency* (JA) (Sosa 2015).

But systematicness brings its challenges. If particular sources of knowledge are each to be brought under a general account, then in developing theory there is a dialectical interplay between the demands of the general theory, and intuitive data about particular cases. I believe that testimonial knowledge – knowledge at second-hand from the spoken or written word of others – presents a particular challenge to general accounts of knowledge. I have argued elsewhere that intuitive data regarding testimony, if accepted, throw doubt on Safety¹ as a necessary condition for

¹Safety, as opposed to Nozick's Sensitivity (Nozick 1981), as a necessary modal reliability condition on knowledge, was first argued for in Sosa (1999). However Sosa's later work develops his 'AAA' account of

knowledge.² I'll give the kinds of case that throw this doubt shortly. The question I address here is: can Sosa's virtue reliabilism avoid the problems for a simple Safety condition on knowledge that testimony presents? In other words, can the problem cases that intuitively are cases where the testimony-receiving subject gains knowledge, although her belief is not Safe, be theorised as satisfying the more complex 'AAA' condition that a belief be accurate, adroit and apt – accurate due to its adroitness – that Sosa's theory proposes as constituting (animal) knowledge?³

I will argue first, that the elegant AAA framework does not determine how it is to be applied to cases, and choices must be made regarding this. Second, that while we can indeed apply the AAA theory to our problem testimony cases to get the intuitively correct results, in doing so we are driven by our intuitions about which cases should count as knowledge. So the AAA framework fails to determine what counts as knowledge. Thus my investigation sounds a common theme regarding reliabilist accounts of knowledge generally: that, even if extensionally correct, they do not provide a reductive account of what knowledge is; since we perforce invoke our intuitions about knowledge in resolving the notorious 'generality problem' of how to individuate reliable processes – or, in the case of the AAA theory, competences and their circumstances of exercise.⁴ This does not invalidate the AAA theory, but diminishes its status, making it less explanatorily significant.

2. The problematic testimony cases

Here are three possible situations regarding a recipient's formation of a belief from taking the speaker's word, accepting as true what she is told. In all of them her resultant belief is not Safe – she might easily have formed a relevantly similar belief that is false.

Safety-Failure Case One, Ignorant Speaker. A person solicits and receives accidentally true testimony that P from a speaker who lacks knowledge and has a true belief regarding P only by chance. This speaker might easily have had a false belief regarding P, and would have misinformed her; she would have believed him in either case.

Safety-Failure Case Two, Patchy Speaker. A person solicits and receives testimony from a speaker on a topic – say, locations of and routes to get to various streets in East Oxford. She asks him where Marston Street is, to which he knows and gives her the correct answer. She might easily have instead asked him about another street, whose whereabouts he would not have known, but still would have answered confidently

knowledge which is more subtle and diverges extensionally from simple Safety. Sosa (1999) defines Safety thus: 'Call a belief by S that P "safe" iff ... as a matter of fact, though perhaps not ... of strict necessity, not easily would S believe that P without it being the case that P.' Explicating Sosa's 'not easily' locution in terms of possible worlds and a nearness metric on them, a belief by S is Safe just if there are no nearby worlds in which S believes it and it is false. More recent accounts have modified this definition, to say that there are no nearby worlds in which S falsely believes P or a relevantly similar false belief. This refinement is needed, for one thing, to cope with belief in necessary truths. 'Relevantly similar' is standardly understood to incorporate relativisation to the method by which the belief is formed. I capitalise the 'S' to mark that this is a technical term with a precise definition, and may not correspond precisely to our ordinary language notion of safety.

²See Fricker (2016).

³Sosa distinguishes a basic kind of knowledge, animal knowledge, from higher kinds: reflective knowledge, where there is reflective endorsement by the subject of her first-level belief as attained with suitable safety; and knowing full well, where her first-level belief is aptly guided to aptness by her meta-level belief. In this discussion I focus on whether a recipient's beliefs from testimony in our problem cases can be portrayed as apt; I will not discuss whether Sosa holds that animal knowledge amounts to knowledge simpliciter – a topic on which his views are nuanced.

⁴See Conee and Feldman (1998).

though incorrectly. In this second case we have a speaker who speaks from knowledge, but accidentally so. He might easily have instead given relevantly similar false testimony. The recipient would believe him in either case. (There are many variants we can spin on this type of case.)

Safety-Failure Case Three, Uniquely Reliable Speaker. A person solicits and receives testimony from an honest and knowledgeable speaker, who is extremely reliable about all the topics she might easily ask him about. But she is in an environment where he is the only reliable informant, and there are many wholly unreliable informants around. It is only by chance that she asked the knowledgeable speaker, and not one of the unreliable ones. She would believe the addressed speaker in either case.

These cases were first cited by me in ‘Unreliable Testimony’.⁵ In that article I maintained that we do not think the testimony recipient (T-recipient) gains knowledge in the first kind of case; but that we do think so, in the second and third kinds of case. I used these claims to argue against a simple Safety condition on knowledge. Instead I argued in favour of a kind of commonsense causal-cum reliabilist theory. I suggested that we have a folk ontology of certain reliable sources of knowledge, and folk theory of how they properly operate; and we think that, when one gains a true belief from such a source operating in its proper fashion, one gains knowledge oneself; even in a situation where one might easily have gained a similar belief from the same source-type (instead operating improperly) that is false.

Testimony as an epistemic source has a unique feature, namely, that the source is herself another human agent, and the information she provides is provided through an intentional act of communication by her. Thus, when testimony operates in its proper fashion, the source herself possesses knowledge.⁶ This enhances the intuition that we get the goods, they are passed on to us – not just true belief, but knowledge – when the source operates normally; even if it might easily have failed to do so, not offering any knowledge to be passed on. My topic today is to consider whether unlike Safety, which fails our testimony test, Sosa’s more complex virtue-reliabilist condition, as formulated in his AAA account of knowledge, can cope with testimony – that is, can deliver the intuitively correct verdicts on our three problem cases.

Our challenge for Sosa’s AAA theory divides into two sub-questions: first, can the AAA theory deliver the result that in Cases Two and Three the recipient does get knowledge; and second, can it at the same time deliver the result that the recipient does not get knowledge in Case One? We will see that Sosa’s theory has the resources to deliver the required negative verdict in Case One; but our investigation will show that the AAA framework on its own cannot deliver the positive verdict needed for Cases Two and Three. How so? It seems first-off that a positive answer to our question can readily be given, as an instance of the more general fact of AAA-theory that a performance can be apt, although it is not Safe. This result is developed about different, but structurally similar cases, by Sosa himself. However we will see that there are underlying complications in applying Sosa’s AAA framework to particular cases: whether a particular attempt comes out as apt or not depends on the description we select of the circumstances in which it is made. In [Section 3](#) I describe the AAA theory, and the contrast between aptness and Safety. In [Section 4](#) I show how a choice always exists regarding any attempt, as to how we describe the circumstances of the agent’s attempt, and what choice we make determines whether or not her attempt is portrayed as an apt exercise of a competence. In the remaining sections I consider what Sosa’s strategy

⁵Fricker (2016).

⁶See Fricker (2015) for an account of how testimony properly functions as an epistemic source that vindicates this claim.

on this matter is, what verdict it gives in our testimony problem cases, and whether a solid rationale for it can be given. My conclusion is that AAA theory can deliver the desired verdict on our testimony cases only when supplemented by our independent intuitions about knowledge.

3. AAA Theory and the contrast between aptness and safety – a quick answer to our question?

Sosa develops a framework that describes performance normativity in general, which applies to any *endeavour*: a skilled activity with a constitutive aim. A performance of the activity is *apt* if and only if it is *accurate* (succeeds in its aim); is *adroit* i.e. manifests the subject's relevant skill or *competence*; and moreover its accuracy is sufficiently due to manifestation of the subject's relevant competence (Sosa KFW p. 4). Believing in the endeavour to attain truth is a specific kind of performance, to which the AAA (accurate, adroit, apt) schema applies. Sosa suggests that apt believing constitutes a basic kind of knowledge – animal knowledge. (A higher kind, reflective knowledge, requires also meta-aptness.) Thus aptness, in Sosa's AAA theory, replaces Safety as the core component in what constitutes knowledge: knowledge is apt belief⁷ (Sosa KFW p. 11).

Our T-recipient's belief in Cases Two and Three is not Safe; yet intuitively, she gains knowledge. So our challenge for the AAA theory is: can the recipient's belief in these cases be portrayed as apt, thus as constituting animal knowledge, despite not being Safe? Sosa's general answer to such challenges is: Yes.⁸ He shows how, as a general matter, a performance can be apt without being Safe;⁹ and argues in several specific cases that a belief (more strictly, a believing in the endeavour to attain truth) is apt though not Safe.¹⁰ A performance is apt just if it was successful, and 'its success manifests a competence seated in the agent' (KFW p. 7). And, crucially, in considering whether this second condition is fulfilled, Sosa has it, we must consider her attempt in the specific circumstances in which it was made. So long as, in those specific circumstances, her relevant competence, duly exercised, suffices for her success (that is – not easily would she fail, when exercising her competence in those specific circumstances), then, Sosa holds, her success is sufficiently due to its manifestation. And, Sosa points out, this can be so, while the success of her attempt¹¹ was not Safe – it was not Safe since she might easily have been in different circumstances in which her attempt

⁷Note: Sosa offers this as a metaphysical analysis, an account of what knowledge is, rather than as a linguistic or conceptual one. See JA Ch. 1.

⁸Sosa does not explicitly consider the kind of testimony cases here proposed. But what he says in relation to Barney and other similar examples is a direct parallel, and so I infer he would propose a similar response. In any case, the AAA theorist has the option to do so – subject to the complications to be explored in the discussion to follow.

⁹See KFW p. 7: 'A performance is apt if its success manifests a competence seated in the agent (in relevantly appropriate conditions). It does not matter how fragile was the competence or its appropriate conditions, when the agent issued the performance. Hence a performance can easily fail to be meta-apt, because the agent handles risk poorly.'

¹⁰See the discussion of how wishful thinking might, but does not interfere with adroit cognitive processing in KFW Ch. 2.3. On KFW p. 25 Sosa writes: 'You see a surface to be red, as you view it in good light. What if the light could easily have been bad? So long as the light is good, I'd say, you can manifest your fine color eyesight in believing the surface to be red. And you can do so even if the light could easily have been red, so that you would have believed the surface to be red, even if it were white, which it might easily have been.'

¹¹Note that to capture the intuitive notion of Safety of an attempt we must gloss it as: not easily would the subject have made a similar attempt that would have failed. Arguably, that very attempt has its particular circumstances essentially, in which case *it* could not easily have failed.

would not have delivered success, and would still in those circumstances have made an attempt – and failed.¹² We can accept this general account of what it takes for a performance's success to be sufficiently due the manifestation by the performer of her relevant competence, and the contrast between aptness and safety that it engenders. But, as will be explored below, there are delicate issues when we come to apply it to cases.

Aptness is one thing, meta-aptness another. A performer's attempt is meta-apt if it is made in circumstances in which it is likely enough to succeed, to make the attempt well-judged,¹³ and it is made in light of her appreciation of this fact. It is fully apt just if it is apt because meta-apt. As Sosa shows, meta-aptness can be present without aptness, and vice versa (see KFW pp. 8–9).

If a performer assesses risk well, and chooses when to make an attempt in light of this, then she will only make an attempt in circumstances in which she is likely to succeed. Hence if a performance is meta-apt, then it is also Safe, since its being meta-apt means that, even if there might easily have been different circumstances, in which her attempt would have failed, she would not have made any attempt in those circumstances. But a performance that is apt, but not meta-apt, may not be Safe: it may be that circumstances might easily have been such that an attempt by her would have failed, and she would still have made such an attempt. So she might easily have failed.

These abstract points are well illustrated with a simple example. Consider two basketball shooters, Anna and Betty. They each have an equal first-order skill in shooting: they can reliably pocket the ball when within a certain zone around the basket – call this their competence-zone. But Anna also chooses her shots wisely, and would not attempt to shoot unless she knew herself to be within her competence-zone. Betty in contrast often shoots (and misses) when she is outside her competence-zone. Consequently, Anna's shots (virtually) all succeed – they are both apt, and meta-apt. Whereas, while some of Betty's shots – those made from within her competence-zone – are apt, others fail, or if they succeed do so only by luck. None of her shots are meta-apt. Anna's shots succeed, and they are Safe because she would not make an attempt in modally nearby circumstances in which she is unlikely to succeed. In contrast, Betty's shots, made with the same level of first-order skill, are not Safe because she would unwisely also attempt a shot in nearby circumstances outside her competence-zone.

This example illustrates how Sosa's AAA framework pans out in a simple application. We see how first-level aptness can be achieved in a performance that is not Safe, and only meta-aptness suffices for Safety. So is it just a one-liner to apply this contrast between aptness and Safety to the particular case where the performance is forming a belief in the endeavour to attain truth, including the case where the belief is formed by understanding what one is told, and believing it?¹⁴ – So that my topic requires only a note, not a full article? The answer I will end up giving, is that we can agree with Sosa that a belief can be apt, though not Safe, and so the AAA framework

¹²See JA p. 72: '[Re basketball player in fact within her shooting-competence zone on court] properly situated as she is *in fact* ... she is thereby disposed to shoot successfully and aptly. Because she is unaware of her threshold, however, she might too easily have shot inaptly. She might so easily have been *improperly* situated and still have shot just the same. [Moreover] even if the lights *might* then easily have dimmed ... so long as in fact they *did not*, her shot could still be apt. Two things are here plausibly compatible: first, she might too easily have shot inaptly (since the lights might so easily have dimmed); but second, she in fact shot aptly ...'

¹³Sosa spells this out in terms of the agent's recognition of the chance of success being 'high enough' and of failure 'low enough', but says little about what determines this.

¹⁴Sosa acknowledges that it is possible to form a belief while not aiming at truth – e.g. through wishful thinking. Only beliefs formed in the endeavour to attain truth are candidates for knowledge.

allows us, if we wish, for instance, to say that Barney, who sees a real barn out of the car window in terrain where there are many fake barns, gains apt belief that constitutes animal knowledge; and to say that a recipient who gains a true belief from a knowing testifier gains belief aptly, even though she might easily have instead consulted one of the many liars around her, and would have acquired a false belief from any one of them. But before getting to this conclusion, as I will spell out, there are complications in applying the AAA framework to cases that must be brought to light. As regards competences, we will find that there are key disanalogies between a competence to succeed in a single action, like pocketing the ball in basketball, versus more complex extended multi-action competences like driving a car; and in particular the logical structure of the competence is more complex and, as we shall see, intrinsically modal, in the case of a discriminatory capacity, such as a perceptual recognitional capacity. These complicating factors make applying the AAA structure to the exercise of a complex multi-action ability like driving, and to perception and other belief-forming competences, much more tricky. My conclusion will be that while one can deploy the AAA structure to vindicate the conclusion that Barney gets apt belief which is animal knowledge, and that our T-recipient in Cases Two and Three also does so, in doing so one makes choices in applying the AAA framework which the framework itself does not entail or justify. If this is right, then Sosa has not given us a reductive account of knowledge, since our intuitions must be deployed in applying the AAA theory.

4. Applying the AAA theory: the need for choices, and Sosa's strategy

A performance is apt just if its success is sufficiently due to manifestation by the performer of her relevant competence. So, in deciding whether in a particular case a subject's performance was apt, we must decide two things: we must decide whether her success was sufficiently due to her manifestation of her relevant competence, as opposed to being due partly to luck (it is clear that this is the contrast Sosa has in mind); but as part of this, we must decide whether she actually manifested a competence at all – whether, in the circumstances of her performance, she possessed an exercisable competence. In what follows I will focus on the first question; but in answering it, the question of whether the subject in various situations possesses an exercisable competence will come up. To address our questions we need to examine more closely the structure of a competence.

A competence is a competence for some action-type ϕ , to ϕ successfully, when one tries (JA p. 96). However, to satisfy the conditional: 'if one were to try to ϕ , one would succeed' is not sufficient for a competence; one must be disposed to try, and to succeed, across a 'sufficient spread of circumstances' (JA p. 97). We can agree with Sosa that this requirement is needed, for the idea of a competence to avoid becoming almost trivial; and also agree that what is a 'sufficient spread' need not be capable of being explicitly defined, being something we have a tacit grasp of.

A complete competence comprises Seat, State and Situation. The core or narrow competence is the Seat. Sosa does not give much detail on the matter, but I take it the Seat of a competence is a specific mechanism of some kind in the subject which is the categorical basis of her disposition to succeed – e.g. when the skill is a visual perceptual one, the relevant cognitive processing module in the visual system in the brain. This Seat constitutes an exercisable competence only when its required State of the subject obtains, and also the required Situation i.e. environmental conditions (call these the *C-environment* of the competence).

The State (i.e. the condition of the subject) for a competence will include conditions like being awake and alert, not drunk or otherwise temporarily physically or

psychologically impaired. The State and Situation will include conditions needed for the core ability to be exercised at all, and others which are required for this to be reliably successful, and as such constitute a competence. For instance, re Anna and Betty: their ball-pocketing core competence can only be exercised when they are in the needed state: awake and physically active and in suitably normal psychological condition; and when they are in the right environmental circumstances – within their competence-zone on the basketball court. Outside this zone, but still on the court, they can in a sense exercise their skill – they can attempt a shot; but it is not strictly a competence outside its required C-environment, since it does not reliably yield success. I will consider that a competence is exercisable only within its C-environment. This accords with Sosa's usage.¹⁵ The following discussion focuses on the complications that arise in applying the AAA structure to cases, as regards the C-environment element in a competence.

Consider again Betty, when she takes a successful shot from within her shooting competence-zone. She is within this zone, so her shot is very likely to succeed. So, it seems immediate that it is apt: successful due to the manifestation of her relevant competence. It is not Safe – she might easily have been outside her competence-zone on the court, and would still have attempted a shot; but she was within it, and so her success is due to her competence, and not at all due to the luck of where she happened to be on the court when she made her attempt. Or so Sosa would have us agree. Not so fast. Betty's complete basketball shooting competence has its three components, and the C-environment component of this, is that she must be attempting her shot from within a certain zone of the court. This is a *type* of circumstance. Equally, when we evaluate her shot as apt or otherwise, we must do so relative to a specification of a *type* of circumstance in which it is made. And, to get the result that it is apt, despite not being Safe, we must specify the type of circumstance in which she takes the shot no less narrowly than that of her competence-zone. If we instead specify the circumstance in which she takes her shot as being simply: in her team's half of the basketball court, then relative to that specification, her success comes out as due in part to luck, not sufficiently due to her manifestation of her competence, since, relative to that type of circumstance, it is (presumably) only luck that she took her shot from the sub-part of her team's half of the court that is her shooting competence-zone.

Now suppose Betty is not such a great basketball player, and there is just one very small specific spot on the court from which she can reliably pocket the ball when she shoots; suppose moreover that she is unaware of where this safety-spot for her is, and is not sensitive to whether she is in it or not. She possesses a shooting-competence relative to that very narrow C-environment, that specific safety spot for her. Now, even if she happens to be in it when she makes an attempt, is it methodologically justified to treat her being in this precise spot as the circumstance of her attempt – rather than some less restrictive, perhaps more intuitively natural specification of her location? Her success will come out as due to manifestation of her relevant competence just if we choose her small safety-spot as the specification of her environmental circumstances

¹⁵A competence is still possessed when the subject is outside its C-environment; it just cannot be exercised. Cf. JA p. 100: 'Clearly one can be a good safe driver, in possession of a competence to drive well and safely even when the nearby roads are all covered with oil. One's status as a good driver is not beholden to the condition of nearby roads.'

Note that, since a competence can only be exercised within its C-environment, this does not leave much scope for one to be exercised, while success is not due to its manifestation. However, anomalies can occur within the C-environment, such as freak compensating gusts of wind for an archer, so this is still possible. (In this case the shot is both accurate and adroit, but not apt – its success is not sufficiently due to the archer's manifestation of her competence.)

when she shoots – but what would justify this choice in applying the AAA theory to her circumstances? In this case one might complain: by putting her precise location into the description of her circumstances, we have made her success artificially come out as due to her manifestation of her competence; but intuitively, it is to a considerable extent due to luck – luck that she happened to be standing exactly where she was, when she took her shot – remember that she takes shots from all over the place, and they mostly miss, since she is mostly outside her competence-zone, this being so restricted. (One's sense that this restrictive specification of Betty's location is artificial is enhanced by the fact that Betty herself is insensitive to whether she is in it, in determining whether to attempt a shot.)

You may or may not find the point I am making compelling in relation to this example. It is used as a simple case, to introduce a general point: the C-environment specified in the specification of a complete competence is a type; and equally, when we consider the circumstances in which a subject makes her attempt, we necessarily are characterising them by means of a type. And, in a case where the subject's performance is not meta-apt, and so is not Safe (i.e. environmental conditions outside the subject's C-environment might easily have obtained, and she would in those conditions have made an attempt and failed),¹⁶ in order to get the result that the subject's success is nonetheless due to her manifestation of her competence, with no significant role for luck, we must select a specification of the circumstances of her attempt that is within the C-environment of her competence. If we instead select a broader specification of the circumstances of her attempt, one that is not within her C-environment, we would instead get the result that her success is in part due to luck as to how precisely she was situated. But this selection of a relatively specific description of the subject's circumstances as she makes her attempt is not compelled by the AAA framework itself, and is a further choice we must make, in applying it to cases. So the AAA framework does not determine which performances, in particular which cases of forming a belief in the endeavour to attain truth, are apt; only further choices we make in applying it do so. We will see shortly how, in regard to our three testimony problem cases, our choice of specification of the subject's circumstances will make a key difference to whether her belief-forming performance is portrayed as apt.

What rationale might justify a particular strategy for applying the AAA framework to cases? First, what strategy does Sosa himself use? On this matter, there is some divergence between his approach in KFW and in JA. I will first spell out and consider the implications and defensibility of the strategy in KFW. In KFW Sosa tells us on p. 7 that: 'A performance is apt if its success manifests a competence seated in the agent (in relevantly appropriate conditions.) It does not matter how fragile was the competence or its appropriate conditions, when the agent issued the performance.' And in his discussion of competences as a species of dispositions in Ch. 2.7 he says that 'dispositions are rarely if ever required to satisfy [a] test of robustness across neighbourhood'; thus, for instance, 'a glass breaks through fragility when dropped on a hard counter, even if all nearby surfaces are cushioned' – that is, its breaking is due to its manifestation of fragility, and not partly due to bad luck in happening to fall onto the one hard surface in the neighbourhood. We get this result if we specify the environment in which the glass's disposition is exercised as: when above a hard surface; we do not, if we specify it as: in an environment with many soft and few hard surfaces. Relative to the latter environment, the glass's breaking is explained only by its fragility

¹⁶A performance's failing to be meta-apt does not quite entail that it is not Safe, since it could be that as it happens there are no nearby circumstances in which an attempt by her would fail, in which case apt selection of when to make an attempt by her is not needed, to avoid risk of failure.

plus the contingency, relative to the specified environment, that it happened to be positioned above a hard surface. The same goes, *mutatis mutandis*, for whether Betty's successful shot taken from her one small safety-spot on the basketball court comes out as apt or not.

These remarks by Sosa suggest that his strategy in KFW is: however limited the subject's competence is, however specific the conditions that must hold, as its C-environment; so long as the *actual* circumstances in which the subject makes her attempt (forms her belief) are included within this, we may, in determining whether her belief is apt, specify the circumstances of her attempt as being the type that define her C-environment; so that her performance comes out as apt – sufficiently due to manifestation of her competence, not due to luck. Thus, in the case of Betty, whose shooting competence is confined to one small spot on the court, we may specify that as her location, assuming she is in fact within it; even though it is a tiny part of the court, and it is mere chance that she is in it, and shoots from within it, rather than from some other location.

5. Applying Sosa's strategy to our testimony cases

What results do we get if we apply this strategy of maximum specificity in description of the subject's circumstances to our testimony problem cases? To address our question we need first to specify the situation and skills of a broadly competent T-recipient in a bit more detail. Sosa's elegant discussions proceed for the most part at a high level of abstraction. But in applying the AAA theory to our testimony cases we need to descend into some more detail. Taking a shot, in hunting or in basketball, is a matter of a single action. In contrast, processes of belief-formation are often a compound of more than one stage – for instance when one forms a belief from perception plus inference (e.g. one sees tiny droppings and nibbled holes in the bananas, and so infers with a sigh that there is a mouse about in the kitchen). In applying AAA theory to such complex belief-forming processes we need to distinguish these different stages, and generally speaking the whole will be apt just if each stage is apt.¹⁷ In the case of a recipient of testimony, we need to distinguish three stages. First there is comprehension of the perceived utterance. Competence in this stage consists in a broad interpretative competence, comprising a narrowly linguistic component, plus an interpretive component needed to resolve the reference of context-sensitive expressions. Its output will be an experience of the perceived utterance as a speech act with specific content and force: the subject literally *hears* the utterance as, say, an assertion that P.¹⁸ This interpretative competence is a capacity to discriminate the content and force of utterances, and it will indeed be a competence if it reliably delivers the correct result – successfully identifies the speech act speakers make, using a suitable linguistic vehicle for this – over a relevant range of possible utterances that are likely to be encountered (say, well-formed sentences of English that are in the repertoire of use in the subject's community).¹⁹ As

¹⁷There are exceptions to this, when a complex competence with many components exhibits a kind of holism – so that overall aptness of a performance may include appropriate compensatory action for small errors at a particular stage. Thus, for instance, a competent golfer will be able to choose a suitable shot to recover, when she lands in a bunker. Thanks to Lisa Miracchi for this point.

¹⁸See Fricker (2003).

¹⁹There is a double generality here: first, a normally competent language-using T-recipient will be able correctly to interpret a range of sentences of her language; but even for a recipient who understands only one sentence of a language, there is a generality in her discriminatory capacity for the force and content made in utterances of this one sentence – she must get it right, for that utterance in contrast with a relevant range of other possible utterances. The idea of a discriminatory capacity is inherently modal.

with the relation between perceptual experience and belief generally, experiencing the utterance as an assertion that P is not the same thing as believing it to be such; and there can be cases where the perception is defeated, so belief is not formed. However in the vast majority of cases the subject will believe the speech act is as perceived, and for present purposes we can neglect the gap between perception of the utterance as a certain speech act, and belief that it is such. Forming a correct and apt belief as to what speech act was made is one stage; accepting the speaker's word and forming belief in what she asserted is a further stage.²⁰

It has been suggested that knowledge from testimony poses a problem for virtue reliabilism, since the credit for the truth of what is believed is due to the epistemic competence of the speaker, not the hearer.²¹ However this objection overlooks that, in order reliably to obtain true beliefs from believing what one is told, the recipient must possess and exercise her interpretative competence, in order to gain knowledge (apt belief) as to what speech act is made.²² Thus there is always a competence exercised by the recipient of testimony, and so some of the credit accrues to her, for the correctness of the belief she acquires through believing what she is told.²³ However her getting of knowledge from believing the speaker requires that the speaker express knowledge; to deal with this, Sosa suggests a plausible extension of his AAA theory, to allow that in the case of knowledge from testimony the competences manifested and explaining success that render the recipient's belief apt include the recipient's, but also include the speaker's competence, in forming the apt belief which she then tells to her audience. We can happily accept this extension of the AAA theory to cover testimony. A T-recipient might uncritically accept whatever she is told, so long as there are not blatant defeaters for it; or she might have her own more or less powerful T-filter, by some means detecting and blocking belief in defective (lying or mistaken) testimony. But even for an uncritical recipient, the need for competence in interpreting utterances means that some credit for her gaining a true belief from believing what she is told accrues to her. In applying the AAA theory to our testimony cases, I will for simplicity assume we are considering an uncritical T-recipient.²⁴ Our T-recipient's testimony-receiving competence comprises her interpretative competence plus her disposition to form belief in what she is told. This will indeed be a competence, i.e. will reliably yield only true beliefs whose truth is due to her manifestation of her competence, and as such are

²⁰These are distinct stages, notwithstanding that very often, in listening to a trusted speaker, one immediately forms belief in what one is told.

²¹See Lackey (2007).

²²It might be objected that, when one is listening to a trusted speaker, one simply 'drinks in' the information she supplies, and does not as it were bother to form a belief as to what speech act is made, instead directly incorporating the asserted content into one's beliefs. I think this can be so; but one must always be in a position to know what speech act was made – this knowledge only a moment's attention away. In my own preferred epistemic framework, this is because one's justification for one's T-belief requires knowledge that it was told to one; in Sosa's AAA framework, such knowledge will be required for reflective knowledge, though possibly not for animal knowledge. An interpretative competence generating apt perception of the speech act is essential, whether or not it must be converted into belief.

²³In the case of an uncritical recipient of testimony, the correctness of the belief expressed in the testimony, which the recipient accepts, is entirely due to the competence of the speaker; however the recipient's interpretative competence is what explains why the belief the recipient acquires is the very same belief expressed by the speaker, and as such is aptly true.

²⁴This simplifying assumption is not essential for the existence of cases like Patchy Speaker and Unique Reliable Speaker: even the most powerful T-filter could not screen out all cases of lying or mistaken testimony, since these can be simply indiscernible from either contextual clues or background knowledge; so such cases – where the recipient is at risk of encountering false testimony she would believe – can arise even for the most critically discerning recipient of testimony.

apt, so long as the speakers whose testimony she receives express knowledge in what they tell to her. That is to say, knowing-speaker is the C-environment of her testimony-receiving competence.

Applied to our testimony cases Two and Three, Patchy Speaker and Unique Reliable Speaker, Sosa's strategy of maximum specificity accordingly yields as a specification of our T-recipient's circumstances: forming belief from accepting the word of a testifier who speaks from knowledge. This specification of the recipient's circumstances yields the result that her belief is apt in both cases; as such, she gains animal knowledge. It does so despite the fact that in Case Two the speaker might easily instead have given similar false testimony, which the recipient would have believed; and in Case Three, the recipient might easily instead have believed another lying speaker's testimony. Under this occasion-specific description of the recipient's circumstances, her belief comes out as apt. Under a less specific description of her circumstances: for Case Two, receiving testimony from someone who frequently lies; or for Case Three, receiving testimony in an environment where there are many liars, it does not do so. So, as regards our motivating question: Can Sosa's AAA theory cope with our problem testimony cases? – the answer as regards Cases Two and Three is: yes it can – that is, it can deliver the intuitively correct verdict that the recipient gains knowledge in both these cases – so long as Sosa's strategy of selecting a maximally specific description of the circumstances of the subject's attempt (in this case: of the T-recipient's environment), one within the C-environment of her competence, however restricted this is, can be defended as having a solid rationale. In the next section I probe some complications and problems for this strategy.

I'll finish this section by dealing with Case One: a T-recipient who acquires a belief from trusting the word of a non-knowing speaker who voices a belief that by mere chance happens to be true. Sosa's theory can readily yield our intuitive verdict on this case, that the recipient does not get knowledge. We explained above how he extends his AAA theory to cover testimony: for aptness of the recipient's belief, the speaker's belief would have to be apt, i.e. be knowledge, and it is not. So the remainder of my discussion concerns Cases Two and Three.

6. Competences for complex skilled activities: a modal C-environment

We have seen that there is a way of applying AAA theory that gets the intuitively correct results regarding our three testimony cases. But does this way of applying it for Cases Two and Three have a solid rationale? In particular, is it legitimate to describe the recipient's circumstances in Cases Two and Three as 'receiving testimony from a knowing subject', when she is in a situation where she might so easily have received false testimony? Surely her close modal environment, not just her brutally actual environment, is relevant to whether her belief is sufficiently due to manifestation of her relevant competence, not due in part to luck? In this section I will consider some more complex competences, attempting to draw out the moral that, for the exercise of some competences, if not all, we must take as the environment in which they are exercised their close modal environment, not simply their brutally actual one: whether the subject has an exercisable competence, and whether her success is due to its manifestation, not in part to luck, turns on the facts about what might easily happen in her actual environment, not merely on what in fact happened, as it turned out.

Consider competence at driving a car. To have such competence, one must be able to cope with the driving situations that one is likely to encounter while driving, as well as others that while not likely, are within the normal parameters of one's driving environment, and so might easily occur. Now consider a person who has a driving competence

for the usual driving situations, but who is unable to steer her car accurately in reverse for any significant distance. Suppose she successfully drives from Oxford to London, via major roads and motorway. No eventuality occurs that obliges her to drive her car in reverse. Is her driving success sufficiently due to her manifestation of her exercised competence – although she is unable to reverse? In this case we can answer affirmatively: though a situation on the M40 could possibly arise, in which she is obliged to reverse the car, this would be highly anomalous. It did not in fact happen, and the possibility it might happen is a remote one that need not be included in what we take to be the circumstances of her driving performance.

But now consider this same driver when she takes a holiday in Devon. Here I must explain that in the county of Devon in England many of the roads are narrow single-track lanes with high banked sides, with passing places where the road briefly widens every five hundred yards or so; if one encounters a car coming towards one from the opposite direction with no passing places between us, the driver who has a passing place nearest behind them must reverse back into it, in order for the two cars to pass by each other and continue their respective journeys. Our reversing-competence-lacking driver sets out on a journey from her holiday home to the beach, which requires driving several miles down such a single-track lane.²⁵ As it happens she encounters no cars travelling in the opposite direction in the lane, and reaches her destination safely. (If a car travelling in the other direction had appeared, she would have panicked and been unable to reverse as required, and her companion would have had to take over the wheel.) Is her success in driving to the beach sufficiently due to her manifestation of her exercised competence, and not due in part to luck that she did not encounter a car travelling in the other direction on the single-track lane? My strong intuition is that in this case we cannot say this: it is part of the general parameters of the driving situation in Devon, that one may need to reverse up a single-track road with passing places on encountering a car travelling in the opposite direction. The fact that as it happens our driver encountered no such other car is a matter of luck, and this lucky fact cannot properly be written into the description of the circumstances of the exercise of her driving competence. The competence of driving is an inherently modal competence: it is a competence to cope successfully at the wheel in a range of circumstances on the road that might easily occur, being included in normal driving situations. Thus the specification both of the C-environment of a driving competence, and of the circumstances of any actual occasion on which someone drives, is a modal specification – of a range of types of event that might easily be encountered by the driver in the situation of her attempt, and which she must be able to deal successfully with, in order to be competent. For driving not in Devon, the need to reverse accurately for a significant distance is an unlikely anomaly, and inability to do so does not destroy driving competence. But to drive in Devon a more powerful competence is required, that includes the ability to reverse some distance up a narrow lane without hitting the sides: the normal parameters of the driving situation are in that respect more demanding, and going with this, the specification of the circumstances in which one drives, when one drives in Devon (with respect to which aptness or its lack is determined) must include the modal fact that an event in which one is obliged to reverse the car might easily happen.

This homely example is given to illustrate and I hope make persuasive a general point: that many complex competences are, as just explained, a matter of being able to act successfully in a range of event-types that might easily, but may not actually on any particular occasion when one exercises the competence, occur. Consequently,

²⁵I recently undertook such a trip – this is no philosopher's fiction, but a real-life example! Let me assure the reader that I am capable, if need arises, of reversing back to a passing-place, though with some sweat.

the specification of the circumstances of any attempt, with respect to which it is assessed as apt or not, must also be modal: must include the events that might easily occur, given the general parameters of the situation, not only those that in fact did so.

This will be true, for instance, of what it takes to be a good (competent) golfer, since a golf round includes a whole series of actions in response to a series of emerging events, and not all the variety of such event-types easily possible given the general parameters of the current golfing situation will be actualised in any particular round. Correlatively, the circumstances with respect to which a golfer's success in a particular round of golf is properly assessed to determine whether it is due to manifestation of the golfer's competence, or in part due to luck, must include all those features which, though they in fact did not occur in her actual situation, might easily have done so in it. Thus if, for instance, the golfer in question can make successful shots only when there is no wind, then if the day was windy, but through luck her strokes were mainly made at moments when the wind lulled, then her success is due in part to this luck. If it was a windless day, on a course where wind is rare, then the possibility of gusts of wind need not be written into the specification of the situation, and her success is in these circumstances apt, due to her manifestation of her golfing competence.

I conclude that for many complex multi-action competences, the specification of the circumstances in which the subject makes her attempt and with respect to which its aptness or otherwise is determined, must include not just what actually happened, but what might easily happen given the general parameters of her situation. Her success will be apt – due to her manifestation of her competence – only if the C-environment of that competence includes all those eventualities that might easily arise in the circumstances of her attempt. (Thus, for instance, where the circumstances of a driver's attempt include that she might have to reverse up a narrow lane, and this is not in the C-environment of her driving competence, her success is not apt, but due in part to luck that an event in which she was required to reverse the car did not arise.)

In fact a more careful consideration of single-action competences shows that this point sometimes holds also for these. Consider again Betty's basketball shot, taken from her small safety-spot on the court. In discussing a similar example, Sosa introduces the idea that the lights might at the moment of her shot have dimmed; he says even so, if the lights did not dim, then her shot was apt – due to her manifestation of shooting-competence, not to luck that they did not dim.²⁶ But, I suggest, this is so only if the lights dimming is an unlikely anomaly. If the basketball court is in a location where power cuts are frequent, and the possibility that the lights may suddenly dim at any moment is a parameter of the situation, something that might easily happen in it, then, I submit, this possibility should be written into the circumstances of Betty's shooting, whose success then emerges as in part due to luck that they did not.

If the moral I have drawn from these cases is accepted, then Sosa's simple strategy of taking a maximally specific description of the subject's actual circumstances, in assessing her attempt for aptness, cannot be maintained without modification. In any case where we apply the AAA framework, we have to decide what are the general parameters of the subject's situation in which she makes her attempt, since these determine what possibilities might easily occur and must be written into the specification of the circumstances of her attempt, with respect to which it is assessed as apt or no. One consequence of this, is that the AAA framework itself does not determine what we must

²⁶JA p. 72: '[discussing basketball player's shot] Suppose moreover that the lights might easily have dimmed just as she was taking her shot. Because of this, she might easily have shot inaptly ... Even if the lights might then easily have dimmed, however, so long as in fact they *did not*, her shot could still be apt.'

say about whether particular attempts, in particular situations, are apt; instead, we must appeal to intuitions we garner from other resources, to decide how to apply the AAA theory. In the case where the attempt is forming a belief in the endeavour to attain truth, the intuition we may need to appeal to, is whether we antecedently regard the belief in question as amounting to knowledge or not. If such recourse is needed, then we sustain a familiar response to reliabilist accounts of knowledge: even if the account is extensionally correct, it fails to give a reductive account of knowledge, since our intuitions about which cases are knowledge are needed to apply the framework of the theory. If this is so regarding Sosa's AAA theory, this does not vitiate the theory, nor rob it of its very considerable interest and significance; but it does impact on its overall status.

Sosa's own discussion of a range of test cases of interest for epistemology, as regards aptness, is subtle and nuanced. In it he shows awareness of the point developed above. On KFW p. 25 he says: '... [suppose that] powerful shifting winds ... already sweep the archery field. Even if the path *happens* to calm ... at any given point just as the arrow approaches, allowing it to reach its target unaffected by any wind, does one's then hitting the target manifest one's competence? The best answer is, I think, that it does. But it is not obvious. Alternatively it was just a lucky shot, since it was so lucky that no wind affected the arrow.' Above, I have sought to make the opposing case on this question. We can however agree with Sosa's verdict (KFW p. 81) that a tennis player's successful shot on a wind-free indoor court manifests her competence, despite the presence physically nearby of outdoor courts swept by strong winds (on which she would not succeed). The situation in which her competence is here exercised is on an indoor court, so the possibility of wind outdoors does not infect it.

On KFW p. 81 Sosa tells us: 'Rarely if ever does a recognised, commonsense disposition require that its triggering conditions would trigger its resultant manifestations not only at the very place and time where the host of the disposition is located, but also throughout some wider neighborhood of places and times.' An instance of this fact is that a glass shatters due to its fragility when dropped on a hard surface, even though all nearby surfaces are cushioned (KFW p. 81). In light of the discussion above, we see that the principle suggested by this statement is insufficiently general. For a simple disposition like fragility, with its one-off manifestation of shattering, the time and location of the shattering suffices as specification of the situation of its exercise, with respect to which it is assessed for aptness. But for a more complex ability like driving, the situation-type with respect to which we must assess an attempt for aptness is not just a place and time, but a modal situation-type, one that includes the things that might easily happen given the general parameters of one's driving situation. Sosa here is in danger of overgeneralising from a simple case. We can retain the point, however, that anomalous bad situations that occur nearby (whether physically or modally), but which are outside the parameters of normal driving situations, do not affect driving aptness. The same goes, *mutatis mutandis*, for other complex competences requiring a modal specification of the circumstances in which they are exercised.

And so we can agree with Sosa's verdict in his discussion of the colour expert: the expert certifies the red colour of an object he views in normal lighting conditions; his verdict is aptly reached, despite the presence nearby of other objects that are white, but bathed in red light (KFW pp. 82, 83). However, there is an important observation to make regarding this and other similar cases Sosa discusses: the verdict that the colour expert's judgement is apt is conditional on treating the nearby presence of white objects bathed in red light as an anomaly, so we can regard normal lighting conditions as one of the parameters of the situation-type with respect to which the expert's colour judgements (apt in that condition) are made. If, in contrast, we were in a world in which

red light intermitted with normal white light continually, then the colour inspector would not have an exercisable competence in that alternative world.

The same goes for Barney: the verdict Sosa gives (JA p. 81), that Barney gains animal knowledge of the presence of a barn, when he sees one, despite the nearby presence of fakes, turns on regarding the presence of these fakes as anomalous. Barney has a perceptual recognitional capacity for barns, it must be maintained, which he exercises, because he can discriminate barns from a relevant range of other types of real building – as against these, he would return the verdict ‘Barn’, when and only when he sees a barn. To regard him as exercising this ability when facing a barn in fake barn country we must regard the modal situation-type he is in as being: confronted with a barn, when he might instead have been confronted with a range of other real buildings from which he discriminates barns. This is the range with respect to which his discriminatory ability is calibrated. If, on the other hand, we think of the presence of fake buildings that are only facades as a normal parameter of Barney’s perceptual situation then, relative to a situation with this modal specification (barn/fake barn), Barney simply does not have a competence – a visual recognitional capacity for barns – at all.

7. Revisiting our testimony cases

Let us revisit our testimony cases Two and Three in the light of these considerations. Our T-recipient is in a situation where she might easily receive false testimony, and would believe it. Is it merely luck that she in fact receives testimony from a knowing speaker? – Should we write these modal possibilities of her receiving false testimony into the description of her situation, so that her getting a true belief comes out as due in part to luck that the speaker as it happens is speaking from knowledge? – Is encountering false testimony in these situations like having to reverse when driving in Devon; or is it like having to reverse when driving not in Devon – possible, but not part of the parameters of the situation?

The AAA framework on its own cannot answer this question. We can go either way; which way we go will depend on our intuitions about what the normal conditions of testimony-reception are, and whether we think the subject is getting (animal) knowledge. The AAA framework permits the answer we intuitively feel is correct; but it does not mandate it, and so it cannot provide any positive support for it.

I mentioned that what Sosa says about competences in JA is somewhat different from what he says in KFW. His additional remarks in JA may help to address the indeterminacy we have found in applying the AAA framework to cases; and may be intended to do so. I find what he says in JA somewhat obscure, so my interpretation may not be accurate. What he says limits which dispositions to succeed when one tries will qualify as competences. We can agree with Sosa’s first point – that there must be a ‘sufficient spread’ of different circumstances such that one would make an attempt, and one’s attempt would succeed, in those circumstances; and we can agree that, as with determining whether behaviour is polite (JA p. 101), we have a grip on what is required in any case, although we cannot provide a definition (JA p. 97). But Sosa goes on to make some further remarks I find puzzling. He says that ‘[d]ispositions and competences are relative to community interest’; in particular, what the relevant State and Situation of a competence are, is fixed by the community’s tacit grasp of these parameters. If he intends that a disposition to succeed when one tries constitutes a competence only when it is recognised as such by the community, who have a concept of it, and tacit grasp of its State and Situation, then I cannot agree with this – I think there are lots of competences, and possession of one does not require that one’s community have a concept of it. But it may be that Sosa is here speaking only of the

‘interesting’ competences (JA p. 101); and it may be that the motive for his discussion, is to point towards the idea that, for these ‘interesting’ competences, what the relevant parameters are of the Situation with respect to which the possessor of the Core of this competence (its Seat) must be disposed to try, and to succeed when she tries, is in part conventionally determined. In discussing driving competence above, I suggested that the modal possibilities that must be included in the circumstances in which the driver makes her attempt, are those that come within the ‘general parameters of her driving situation’. And there is certainly room for the idea that, to an extent, what these are is determined by community attitudes to driving and its circumstances, not solely by the natural facts about what is and is not able or likely to happen in the course of a car journey.

If this is Sosa’s drift, then this resource – the community’s conception of what is the C-environment of a recognised competence – may help to resolve the question of how we apply the AAA framework to cases; it may resolve what possible eventualities are taken to be within the general parameters of the situation of an attempt, and what other possible events would be anomalous, outside the normal parameters of the situation.

How could this help determine aptness or its lack in our testimony cases Two and Three? It would help, if the community’s tacit grasp has it that to be a competent receiver of testimony, it is enough that one can aptly interpret the heard utterance, and form belief in it; one need not be able to somehow discriminate between knowledgeable versus false testimony, because givers of false testimony, though entirely possible, are anomalous, not one of the parameters of the general situation of receiving testimony. But to hold this just is to hold that our uncritical T-recipient in Case Two and Case Three does indeed form belief aptly, gaining (animal) knowledge. So – to repeat the earlier charge: while Sosa’s AAA framework achieves much of great theoretical significance, for one thing explaining the normativity and value of knowledge (apt belief) as an instance of the broader performance normativity of endeavours in general, it does not achieve that elusive epistemologists’ Holy Grail, a fully reductive account of propositional knowledge.

References

- Conee E. and Feldman R.** (1998). ‘The Generality Problem for Reliabilism.’ *Philosophical Studies* **89**, 1–29.
- Fricker E.** (2003). ‘Understanding and Knowledge of What is Said.’ In A. Barber (ed.), *Epistemology of Language*, pp. 325–66. Oxford: Oxford University Press.
- Fricker E.** (2015). ‘How to Make Invidious Distinctions Amongst Reliable Testifiers.’ *Episteme* **12**, 173–202.
- Fricker E.** (2016). ‘Unreliable Testimony.’ In H. Kornblith and B. McLaughlin (eds), *Alvin Goldman and his Critics*. Oxford: Blackwell.
- Lackey J.** (2007). ‘Why We Don’t Deserve Credit for Everything we Know.’ *Synthese* **158**, 345–61.
- Nozick R.** (1981). *Philosophical Explanations*. Oxford: Clarendon Press.
- Sosa E.** (1999). ‘How to Beat Opposition to Moore.’ *Philosophical Perspectives* **13**, 141–53.
- Sosa E.** (2011). *Knowing Full Well*. Princeton, NJ: Princeton University Press.
- Sosa E.** (2015). *Judgement and Agency*. Oxford: Oxford University Press.

Elizabeth Fricker was for many years a Tutorial Fellow at Magdalen College Oxford, and University Lecturer in Philosophy at Oxford. In September 2017 she retired from that position, and is now a Visiting Professor at the University of Notre Dame, Indiana, teaching there in the spring semesters. Her main research interest is in testimony, on which she has published very many articles. She is currently preparing a book that brings this work together for publication with Oxford University Press.

Cite this article: Fricker E (2019). Knowing Full Well from Testimony? *Episteme* **16**, 369–384. <https://doi.org/10.1017/epi.2019.31>