A Simpler Model of Judgment: On Sosa's *Epistemic Explanations*

In *Epistemic Explanations*, Sosa continues to defend a model of judgment he has long endorsed. On this complex model of judgment, judgment aims not only at correctness but also at aptness of a kind of alethic affirmation. He offers three arguments for the claim that we need this model of judgment instead of a simpler model, according to which judgment aims only at correctness. The first argument cites the need to exclude knowledge-spoiling luck from apt judgment. The second argument uses the complex model to distinguish judgment from mere guessing. The third argument involves the assessment of suspension of judgment as a performance. This paper shows why none of these arguments succeeds, and so recommends adopting the simpler model of judgment.

In *Epistemic Explanations*, Ernie Sosa refines the powerful and original theory of epistemic normativity he has long defended.¹ His **telic virtue epistemology** (TVE) centers around a kind of doxastically committal occurrence—a **judgment**—that can form or modify a lasting epistemic commitment like a belief.² Sosa's innovation is his treatment of judgment as a **performance**: an action with a specific aim. The epistemic credit a judgment deserves then depends on its quality as a performance directed towards that specific aim.

What kind of aim must this be? According to Sosa, the relevant aim of judgment is an intention to answer some question (say, *whether p*) not only correctly, but also **aptly**, through your own competence.³ It's only this kind of occurrence which is fully attributable to you as the agent, and thus creditable to you. But this is a remarkably sophisticated intention to have. Requiring such an intention for every judgment risks over-intellectualizing what it really takes to make a judgment.

Sosa's treatment of judgments as performances is powerful in many ways. He has

¹ Sosa (2021). I use the abbreviation "*EE*" in what follows to cite pages from this book, and *J*&A to cite pages from his (2015) *Judgment & Agency*.

² Sosa is also concerned to explain the knowledge status of dispositional beliefs and even some noncredal states, but here I will set these aside.

³ Compare his "Knowing Full Well" (2009).

extracted significant insight from this approach to epistemic normativity. But his approach does not require his specific treatment of judgment's aim. In fact, the view can accomplish all it accomplishes now even with a simpler picture of judgment, on which it aims not at aptness but rather just at correctness in answering a question.

To show this, I'll first summarize the reasons Sosa favors this complex model of judgment. He has argued that we need this complex model

- (a) to rule out certain forms of knowledge-spoiling luck in apt judgment;
- (b) to differentiate judgment from mere guessing; and
- (c) to make sense of the performance normativity of suspending judgment.

I'll take on each of these arguments in turn. He himself recognizes in this latest book that this model doesn't help with certain forms of luck that matter to (a); I add that it is not needed to protect apt judgment from other forms of luck either. On (b): the complex model of judgment can't properly differentiate judgment from guessing. As for (c): we can explain suspension and its appropriateness in terms familiar to performance normativity with a simpler model of judgment.

While this paper is critical of the details of Sosa's proposal concerning judgment, it is also deeply sympathetic to the approach behind his telic virtue epistemology. The alternative model of judgment presented here shares a core motivation with TVE: the motivation to explain the epistemic credit we give judgments in terms of their performance quality.

Let me start by sharpening our motivating question: what *is* this complex model of judgment, and what is the alternative?

1. Two models of judgment

One of Sosa's great philosophical achievements is to show why treating judgments as *performances* helps us to understand the norms that apply to them. He explains the epistemic normativity of judgments in terms of familiar features of intentional action.⁴ I'll summarize his way of thinking about this normativity here.

First, consider the role of an agent's own **aim**, here embodied in an *intention*. Whenever you φ intentionally, or even just *try* to φ intentionally, you have a certain aim in doing so—the aim to φ . Your intention to φ sets a substantive standard on your success. When you try to get the right answer to a question—as you do in judgment—your aim is to get that right answer, and that makes sense of the standards we apply to judgment. If your judgment is such an attempt, it's a success only if it is accurate, precisely because you're aiming at the *right* answer to the question. Thus, your aim as the agent explains one way that judgment can be good or bad as a performance: it can meet or fail to meet the standard your intention sets for it.

Second, any performance can also be assessed for the competence the agent manifests in that performance. Even when you do φ in aiming to φ , we can assess the *way* you managed to do that, and indeed question whether success really is

⁴ For explicit summary of the importance of intentional action, see J&A Chapter 7.

creditable to you. In the best situation, you manifest **skill** in your performance, you yourself are in good **shape** to exercise this skill, and the world cooperates with a favorable **situation** of its own. Combining such skill, shape, and situation constitutes **complete competence** in your performance.⁵ Actually manifesting competence in trying to φ is to φ **adroitly**; this is not to guarantee success in your attempt to φ , but it does deserve credit. In the case of an attempt to judge, various forms of competence correspond to various forms of justification.⁶

Finally, we can assess whether some performance is successful *because* adroit. We know from extensive discussions of deviant causal chains in the philosophy of action that you can manifest competence in some attempt to φ , and you can indeed φ in that attempt to φ , without thereby *intentionally* (and so creditably) φ -ing: your success might be due to luck instead. The very best version of a performance is one that is successful, adroit, and successful *because* adroit, i.e. an **apt** performance.

There's yet another standard that Sosa uses to assess performances: the standard of **full aptness**. This is "where the agent aims not just at accuracy [or more generally, success at some embedded aim] but at aptness, and succeeds through competence in this more complex endeavor" (*EE* 20). It's the role of full aptness in judgment in particular that is at issue in our disagreement.

To clarify the disagreement, let's apply all this to an example of judgment. Consider the aim of getting it right on some question, say, whether it's raining. If you act on an intention to correctly answer the question whether it's raining, you might succeed by coming to think it is raining. If you do that through your competence, you aptly attain your aim of getting it right on the question of whether it's raining. This would be an apt performance, but it wouldn't qualify as a *fully* apt performance unless it is also part of your *aim* to do all of that aptly. To get there, you'd need an intention to get it right aptly on the question of whether it's raining the relevant question, but also by exercising your **second-order competence** in guiding your first-order performance towards aptness. Only if you succeed at getting it right *aptly* on the question of whether it's raining—and only if you do that whole more complicated thing aptly, through the exercise of your second-order competence—does your performance count as fully apt.

Sosa usually introduces the importance of full aptness in epistemology by pointing out that an apt but not *fully* apt performance still involves an "element of relevant luck in its success."⁷ This element of luck limits the extent to which we can credit success—such as success in answering a question correctly—to the agent. In the case of judgment, aptly but not *fully aptly* answering a question correctly could not count as gaining knowledge. At least, it could not count as gaining knowledge of a particularly important kind, which Sosa has called "reflective knowledge" or

⁵ "Skill," "shape," "situation," and "competence" are all his terms. *EE* 192ff.

⁶ Just as adroit performance doesn't guarantee success, justification doesn't guarantee accuracy. *EE* 200ff.

⁷ EE 21.

"knowing full well."⁸

In accordance with the importance of full aptness in epistemology, Sosa has defined judgment in terms of the more complex aim you would need to have for your judgment even to qualify for full aptness. He says judgment is an "endeavor (attempt) to get it right *aptly* by alethically affirming that p," where **alethic affirmation** is an "endeavor (attempt) to get it right by affirming that p," and an **affirmation** that p is like saying that p, either out loud or silently to yourself.⁹ In other words, your intention itself explicitly demands not just a correct affirmed answer to a question, but one that is reached aptly. To execute *this* intention aptly, you must exercise a second-order competence to make these kinds of correct alethic affirmations aptly. A fully apt judgment will be one made through the successful exercise of that second-order competence (as well as first-order competence).

This is the **complex model** of judgment that Sosa offers. I'll call the relevant kind of epistemic action "judgment-c" to highlight the complexity of the intention involved—the content of which is in italics here:

Judgment-c =_{def} an affirmation that p made as an attempt to answer Qp correctly and aptly

where Qp is a question that has p as one of its answers.

The clear foil to this account is a less complex account of judgment. On the **simple model** of judgment, the performance on which we should focus is judgments, which executes a simpler intention (again, here in italics):

Judgment-s =_{def} an affirmation that *p* made as an attempt to answer Qp correctly

In what sense are the simple and complex models proposals about the same thing, judgment? What *is* judgment, such that you can have these divergent models of what's involved in it? I'm thinking of judgment, pretheoretically, as the mental act that is the analogue of the mental state of belief. It involves a doxastic commitment to something's being true, and it deserves epistemic credit or discredit depending on how it is done well. Additionally, in the framework of TVE, we can say that judgment is that mental act whose features as a performance explain epistemic credit in its many shades and varieties.

Sosa and I agree that judgment should be modeled as a special kind of intentional action. What we disagree about is the intention on which you act when you judge. Here's another way of expressing our disagreement: I claim that we can explain all that Sosa wants to explain about epistemic normativity with just what he calls "alethic affirmation," i.e. judgment-s. The distinction he draws between alethic affirmation and judgment-c turns out not to be necessary for the purposes to which he puts it.

That is what I will argue in the remainder of this paper. But first, it's worth noting

⁸ See also A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I (2007) and "Knowing Full Well" (2009).

⁹ EE 24.

why it matters whether we endorse a complex or a simple model of judgment.

In the abstract, there are reasons to prefer a simpler model of judgment over a complex one.

First and foremost is the fact that judgment-c requires us to attribute very sophisticated conceptual powers to any genuine judger. Only someone with the concept APT could judge-c. That's because judgment is a matter of doing something with a specific kind of intention, as Sosa and I agree. An intention is an explicit, cognitive, person-level representation of the agent: to intend is to have something in mind to be done.¹⁰ You can only intend something which you can conceptualize, just as having a belief demands you have the concepts in which it's framed. Sosa builds into the intention any judger must have that it's an intention to do something *aptly*, and so anyone who engages in any judgment at all must have the concept rather than an everyday one shared by people going about their everyday lives.

Second, anyone aiming at aptness *as such* should have an embedded understanding of what would count as apt in some particular context of judgment, but this is often very difficult to discern. Even if we treat the aim of aptness as implicit, as Sosa sometimes suggests we must, attribution of this aim to a thinker still requires attribution of such sophisticated implicit understanding.¹¹

We must have good reason to prefer the complex model of judgment, then, if we are to choose it over the simpler model of judgment that only demands that thinkers have the concept CORRECT. But we don't have good reason to prefer the complex model of judgment. To show why not, I'll just show that the three kinds of reasons Sosa offers aren't dispositive.¹² First, he takes it that judgment-c is necessary to rule out a kind of knowledge-spoiling luck. Second, he thinks only judgment-c is properly distinguished from guessing, and judgment-s is not. Third, he takes the more complex intention that judgment-c executes to be essential in explaining the performance normativity of *suspending* judgment.

I'll argue against each of these points in turn, starting with the point about luck.

2. Luck

Sosa has long claimed that fully apt judgment-c rules out an important kind of luck that might afflict apt judgment-s. Here I'll argue that considering this kind of luck doesn't end up providing reason to favor the complex model of judgment.

¹⁰ An aim that's an *intention* is special among aims. Your digestive system can aim to break down food without having any concepts at all, but that's because its aim isn't a matter of an intention. Sosa is sensitive to the distinction. See, e.g., *J*&A p.19: "Perception involves functional, teleological aimings, through the teleology of our perceptual systems. Intentional action involves aimings that are full-fledged intentions. Knowledge divides into two sides: a functional perception-like side, and a judgmental action-like side." Thanks to an anonymous referee for pushing this point. ¹¹ See e.g. *J*&A 84.

¹² Sosa has also used the complex model of judgment to help explain Descartes's *cogito*, but he has also treated this as a "bonus" of the model rather than a motivating reason to accept it ($J \otimes A$ 85). I'll set this aside in this paper.

2.1. Secure knowledge

In earlier work, notably his 2015 Judgment & Agency, Sosa motivated the need for a complex model of judgment partly in response to worries about knowledge-spoiling luck. Consider Barney, who looks at one real barn in a neighborhood full of fake barns and judges of that one real barn that it's a real barn. If his judgment is judgment-s, it's an attempt just to answer a question *correctly*, and as that kind of performance it is apt. However, it would not be *fully* apt as a judgment-c, and would not amount to **reflective knowledge** or **knowledge full well**; for that, "Barney must know that if in his conditions he affirmed that he faces a barn, not easily would he be wrong."¹³ He doesn't know this fact in his context, so his judgment is not fully apt, and he doesn't know full well that what he is looking at is indeed a barn.

This was Sosa's view in *Judgment & Agency*. In that book, Sosa argued that fully apt judgment-c excludes the kind of knowledge-spoiling luck that afflicts Barney's belief, as well as other beliefs like those of Norman the clairvoyant, Truetemp, blindsighters, and chicken sexers.¹⁴ It is guidance by second-order competence that matters here: "they fall short through lack *even* of an apt presupposition—an apt *implicit* awareness—that their relevant first-order affirmations are and would be apt ... This second-order awareness must also *guide one* to the relevant aptness on the first order."¹⁵ More generally, Sosa claimed, in order to avoid epistemic negligence in gathering or assessing evidence, and to ensure sensitivity to overriding and undermining defeaters, judgment must be fully apt judgment-c in order to count as reflective knowledge.¹⁶

In *Epistemic Explanations*, Sosa substantively revises this view. He recognizes that there are a whole host of conditions that must obtain for any performance to be apt which you need *not* know about in order to perform aptly. You must be able to take various "background conditions" for granted in your performance, whether that's a purely bodily or a relevantly epistemic performance, even when those conditions don't hold safely.¹⁷ To do this is to make a "default assumption" that is licensed by the domain of performance in which you participate.¹⁸ For this reason, Sosa admits that Barney and others actually *can* have knowledge full well, because judgment-c in their situations can be fully apt.¹⁹

Sosa then introduces another level of knowledge to explain what Barney actually lacks. He calls this "secure knowledge." It is distinguished from other reflective knowledge by the fact that "you are safe from the following fate: *losing* your pertinent complete SSS competence to so judge while at the same time *retaining* a disposition to make or suspend judgments when you 'inquire' into that question *even absent any*

¹³ J&A 79.

¹⁴ J&A 77ff.

¹⁵ *J&A* 81 and note 22.

¹⁶ J&A 77-88.

¹⁷ EE 124-6

¹⁸ EE 128-132

¹⁹ EE 139-141, 146-7, 168-70. Quotation from 170.

*such competence.*²⁰ This, he says, is a higher grade of knowledge even than reflective knowledge. It is "higher" along a hierarchy of forms of knowledge that tracks "categories pertinent to degree of attributability/responsibility, and opposed to metaphysical luck.²¹

This important adjustment to Sosa's account constitutes an equally important concession concerning judgment. By introducing an additional condition on judgment to exclude the relevant luck, Sosa also concedes that the notion of full aptness does *not* do some of the work he introduced it to do.²² This undermines the original motivation for Sosa's complex model of judgment.

Note that the demand for security, in Sosa's sense, is a very different demand than the demand for *full* aptness of judgment (i.e. aptness of *judgment-c*) over mere aptness of judgment-s. It introduces an externalist counterfactual condition on this high grade of knowledge. Even if we concede that security in Sosa's sense still affects credit to the agent, we should still question the extent to which it depends on fully apt knowledge—and thus on apt judgment-c—at all.²³ The notion of security might instead be introduced as a condition on knowledge constituted by simpler judgments aimed at correctness instead of aptness—i.e. knowledge constituted by judgment-s. Technically, since Sosa defines security in terms of his own notion of judgment, which refers to judgment-c, such security could not apply even to apt judgment-s. But apt judgment-s could still enjoy "security-s," which is precisely analogous to Sosa's security for judgment-c. Whether or not judgment-s is secure-s is just orthogonal to whether the aptness of such judgment-s is guided by any second-order competence, or whether that judgment-s is performed as part of a more complex attempt that could constitute judgment-c.

This suggests that one of the main reasons Sosa originally transitioned from a simple to a complex model of judgment isn't actually reason to make that transition at all. There is a kind of epistemic luck which Sosa now recognizes just can't be excluded by demanding guidance by second-order competence. But if that luck can't be excluded in that way, we might not want to make that otherwise costly move from a simple to a complex model of judgment.

Sosa does still think there are *other* forms of credit-reducing luck that would afflict apt judgment-s but could *not* afflict apt judgment-c. It might be because he wants to rule out these other forms of luck that he only considers judgment-c to be a candidate for knowledge that may or may not be secure. Let's return, then, to those more basic reasons he cites for preferring the complex model of judgment over the simpler one.

2.2. Quality of choice

A second reason Sosa favors the complex model of judgment concerns the choice of *when* (or *whether*) to make an attempt at all.

²⁰ EE 170-1.

²¹ EE 180.

²² J&A 77-88.

²³ See EE 185-7 summing up Sosa's views on security, credit, and grades of knowledge.

We can assess a choice to make an attempt just as much as we can assess the attempts themselves. An athlete might rightly be faulted by their coach for taking a risky shot on goal when a better alternative was available. Similarly, a detective who makes a premature judgment about who has committed a murder might be faulted by the police chief when more evidence should have been considered before making any such judgment at all. The context of choice does indeed seem to affect the performance-relevant assessments of the attempts themselves; the shot is worse in some sense for being poorly chosen, and the judgment leading to conviction is worse in some sense for its being made in a risky way. On these points, Sosa and I agree.

We disagree about whether these points should incline us towards a complex model of judgment. Sosa thinks that they should, apparently because assessing an agent's performance in terms of her *choosing* whether or not to make an attempt at all could indicate that she is thinking about the aptness of a performance she may or may not attempt. But I don't think these points about assessment of choice should incline us in this direction. There are two problems with using these observations as reason to accept a complex model of judgment.

The first problem is this: there *isn't* always a relevant choice to be assessed, either in the case of bodily performances or in the case of judgments. An intentional action might not always be initiated in the context of a real choice between alternatives. (Consider coercion or habitual action.) What's more, some situations so obviously demand *some* attempt or other that even if the choice is made, it's not obviously negatively assessable according to the standards of risk involved in that attempt.²⁴ A shot on goal taken in the last seconds of a match is better than no shot at all.

Sosa can agree with all of this. But he might not have seen that it undercuts the way he uses these observations about choice of attempt to support the complex model of judgment. The fact that just *some* judgments can be assessed for being well or poorly chosen doesn't obviously support the claim that *all* judgments must be performed with an eye to their aptness. What's more, since some judgments positively can't be assessed for how well they are chosen—since not chosen at all—it seems unlikely that reflecting on such choice will lead us to the complex model.

The second problem is this: the fact that we can assess the choice to attempt to φ does *not* demand that we then see that attempt to φ itself as directed at *apt* φ -ing. Once the choice of whether or not to try to φ is complete, and the agent begins in earnest to try, the question of aptness might be left entirely to the side. In such a case, the agent doesn't guide what they do with any idea of aptness in mind, even implicitly; she may just focus on the first-order goal of φ -ing. But that is perfectly consistent with assessing the choice this agent made about to attempt φ -ing. We must distinguish choice about *whether* to try from the aim of the attempt itself.

2.3. First-order aptness and luck

Sosa often claims that an attempt at correct alethic affirmation made only on the first

²⁴ Sosa recognizes this much when he contrasts an archer in competition, who simply has to make a shot to participate, with an archer hunting in the woods.

order—in other words, a judgment-s—simply could not manifest sufficient competence to be apt *in its own right*. Consider the following passage:

Unless aimed at least in part at avoiding inaptness, deliberation [on some question] and its outcome would not be wholly competent. Unless one managed to avoid inaptness well enough, one's pursuit of aptness would fall short, so that any aptness one might attain would manifest insufficient competence. Unless one aims to affirm alethically *only if* one would do so aptly, and one properly guides oneself to do so, one's attainment of aptness is relevantly lucky. It is *insufficiently* owed to competence.²⁵

This is a difficult passage that can be read in several ways. But still, on one reasonable reading, Sosa here claims that the *first-order* attempt wouldn't be competent, and so couldn't be apt, without *second-order* guidance. Then aptness of the first-order attempt would require an attempt *at* aptness on the second order as well. But this cannot be true in general, as it would quickly generate an infinite vicious regress: aptness on the second-order should also demand an attempt at aptness on the *third* order, and so on and so forth.²⁶

To see another issue with this line of reasoning, consider the following claim about apt and fully apt performances more broadly:

If an attempt succeeds aptly without being fully apt, there is an element of relevant luck in its success. Its aptness is not secured through the guidance of the agent's second-order competence. It is thus lucky that the agent succeeds aptly.²⁷

What precisely isn't enough in this case—and what is it not enough for?

In the case discussed here, as in the kind of case mentioned at the end of the last passage, it can't be that the agent doesn't have or manifest enough competence to achieve the first-order aim, or even that the first-order aim isn't sufficiently explained by the exercise of that competence. These are cases in which that first-order achievement *is* apt, by hypothesis. So it can't be that the luck involved spoils apt performance on the first order. It also can't be that the lack of guidance by second-order competence ruins first-order competence required for aptness on the first order. But then it seems like the insufficiency involved in either just is insufficiency for aptness at the second-order. Then we would read Sosa as saying that too much luck in the *second-order* attempt spoils aptness at the second-order—which is simply true by definition.²⁸ That can't give us further reason to choose a complex model of

²⁵ EE 52.

²⁶ In personal conversation, Sosa has responded to this point by claiming that the demands made of any judger must be reasonable, and not too demanding, and so the regress doesn't get going. But if cognitive difficulty is to be considered in this way, we might as well go with a simple model of judgment anyway, since expecting all judgers to have the concept APT is demanding in itself. ²⁷ *EE* 21.

 $^{^{28}}$ The same goes for the point that aptness at the first order isn't enough to meet the *aim* of

judgment over a simple model of judgment. What is missing is a reason to care explicitly about second-order competence, above and beyond first-order competence and its exercise in successful, adroit, and *apt* first-order performance.

Sosa sometimes does cite specific ways in which second-order competence seems to matter to apt judgment. In particular, he thinks that second-order guidance is required to avoid epistemic negligence and recklessness in treating your own evidence.²⁹ To make an apt judgment, he writes, you must be guided towards your judgment not only by the "contents of your body of relevant total evidence" but also by the fact that such evidence is all of the evidence you have in your possession *and* the fact that such evidence is enough to justify judgment over suspension.³⁰

Sosa says this guidance demands a move from the first- to the second-order level: "Note how your thought must here ascend to the second order, by bringing within its scope (at least implicitly) considerations *about* your evidence, including its being *in your possession*, and *being on balance sufficiently extensive and strong*," Sosa writes.³¹ If you didn't base your judgment on these further facts as well, even if your judgments were apt at the first order, such "*aptness* is then attained by luck rather than competence"—luck, I take it, of the same kind that more generally afflicts judgments and does not afflict apt judgment-c.³² This is meant to give us further reason to think that the aptness required for apt judgment-c is epistemically important above and beyond that required for apt judgment-s, and that gives us reason to think that judgment is best modeled in the complex way Sosa favors.

In reply, I'll first grant that implicitly basing your judgment on these assessments of your evidence is important for apt judgment. The question is now whether this requires us to capture the aptness of judgment as the aptness of judgment-c instead of judgment-s. It does not.

Consider a situation in which you were engaged in an attempt at judgment-s, i.e. an attempt to answer some particular question *correctly*. In determining whether an answer to a question is correct, you can be thoroughly engaged in the first-order question of what is correct, without any attention to the aptness of what you're doing *as such* (as would be involved in an attempt at judgment-c). Focusing on what the *right* answer to the question is can just as well involve an implicit assessment of your evidence—and an implicit assessment is all Sosa thinks is needed, after all. This implicit assessment can take place when you consider reasons for and against certain answers to the relevant question directly; second-order thought need not occur.

To see this point more forcefully, consider the difference between (i) thinking of something as *evidence for a certain answer to a question* and (ii) thinking about *your answering the question*. Even if you were to *explicitly* assess your evidence *as such* in the course of trying to answer a question correctly—which Sosa doesn't require—even that

²⁹ EE 59-64.

³⁰ *EE* 60. Sosa makes these claims in part to argue that a purely evidentialist approach to epistemological justification would not be able to handle suspension of judgment effectively. See SIV for more on suspension of judgment itself.

³¹ EE 62.

³² Quotation from EE 63.

would not involve ascent to the second order in the sense that matters to the difference between judgment-s and judgment-c.

A final point here: Sosa sometimes speaks as though conditions of aptness on a certain kind of endeavor make that endeavor what it is. But we must be careful not to confuse conditions on what it is to φ *aptly* with conditions on what it is to try to φ *at all*. Even granting that these kinds of evidence-assessment are required for *apt* judgment, then, would not imply that evidence assessment is required in *any* attempt at judgment.

2.4. Summary

In this section I considered how Sosa uses worries about knowledge-spoiling luck to motivate the complex model of judgment. I pointed out that Sosa's amendments to his discussion of secure knowledge remove one crucial previous motivation for the complex model. Then I looked at two further lines of argument for the complex model of judgment that try to shield apt judgment from knowledge-spoiling luck. Against the first, I argued that you don't always choose whether to attempt a judgment, and that even when you do, it need not affect the aim with which you judge. Against the second, I argued that there is in general no reason to think aptness requires full aptness unless we have some independent reason to care about full aptness. Sosa's considerations about evidence assessment also do not rationalize accepting the complex model of judgment over the simpler one.

I suspect that the more fundamental reason Sosa remains uncomfortable with judgment-s is not because judgment-s can't be properly competent on its own, but rather because judgment-s doesn't even really seem like judgment to him at all; some judgment-s, he claims, is mere guessing.

3. Guessing

It's important that the type of performance so central to TVE is not the same as *guessing* the answer to a question. Sosa's complex model of judgment is motivated partly by the need to distinguish genuine judgments from guesses.³³

What is a guess? Sosa introduces guesses with three types of examples. You might give a **sheer guess** at the answer to a question when you randomly select some answer to assert on a game show, to give yourself at least some chance of winning a prize.³⁴ You might offer a more intelligently guided guess at the identity of the blurry letters low down on an optometrist's eye chart; I'll call this a "**perceptual guess**."³⁵ Or you might, as an "insightful theoretician," intelligently navigate your way through some domain of uncertainty to offer an **educated guess** on some matter of serious importance.³⁶ The methods of these guesses are very different: there's random selection of an answer for sheer guesses; imperfectly and unconfidently perceptually

³³ J&A 85.

³⁴ EE 54.

³⁵ J&A 74ff; EE 144 note 2.

³⁶ EE 162-3.

guided answers for perceptual guesses; and carefully and intelligently guided answers for educated guesses. But all are meant not to be judgments.

Here's how this might put pressure on a simple model of judgment. Judgment-s is just an affirmation made as an attempt to answer some question correctly. Any of the cases above could constitute some such attempt, and indeed could answer the relevant question correctly; a guess can be right, after all, even if only by some degree of luck. But these are *not* meant to be cases of judgment. Thus it seems that something over and above judgment-s is required for judgment. Sosa introduces the complexity of the intention involved in judgment-c partly to fill this out.

The key difference between judgment-s and judgment-c, however, can't properly distinguish between judgment and guessing.

3.1. Aiming at aptness

Initially it seems that aiming at aptness, so essential to judgment-c, is incompatible with guessing. We might say that the guesser does aim at a correct answer, but *doesn't* aim at aptness in achieving that. It's the latter that makes the difference between guesser and judger. That would make a neat distinction between guessing and judging and provide a clear rationale for the complex model of judgment over the simple model of judgment.

But things aren't so straightforward. Sosa recognizes that an educated guess can aim at aptness, as can some perceptual guesses. At least, these guessers "are aiming for some degree of aptness, some degree lying above the nil degree of a sheer guess, even if below the higher degree required for expert or even ordinary knowledge."³⁷ This concession complicates things, and undercuts this first rationale for the complex model of judgment.

3.2. Confidence

Sosa does try to salvage the account of guesses in terms of *confidence*. He claims that you need sufficient *confidence* both at the first order (in whether the answer to the question really is *p*) and at the second order (confidence in the aptness of how you're answering that question) to count as judging at all rather than guessing.³⁸ Such confidence is a requirement for *competence* at either level, and competence—he claims—is a requirement for judgment of the relevant kind.³⁹ This is meant to exclude guesses from being instances of judgment, since in each case—sheer guessing, perceptual guessing, or even educated guessing—some relevant competence-required

³⁷ EE 163. In earlier work Sosa even conceded more simply that the perceptual guesser could aim at aptness simpliciter, and this makes it seem like a guesser could have and act upon the whole complex intention that is meant to characterize *judgment* in contrast to guessing; see *J*&A 75, especially note 14.

³⁸ J&A 75 n14, EE 79ff.

³⁹ In *Epistemic Explanations* Sosa newly introduces further ways of assessing epistemic performances in terms of the role of confidence, which can be aptly or inaptly proportioned to the degree of your evidence. Your competence in judging at all partly depends on your confidence in apportioning confidence. *EE* 80ff.

confidence is lacking, at least at the second order.

This doesn't obviously support the complex model of judgment, because confidence—at either the first or second order—is orthogonal to whether you're aiming at aptness. In a judgment-s, when you are trying to get the right answer on some question without thereby *aiming* at aptness, you might be very confident both that you're getting the right answer and that your way of getting the right answer is apt. Confidence about aptness is one thing; *aiming* at aptness is another.

Of course, judgment-s doesn't demand confidence about aptness at all, since it doesn't even require any thought at the second order. But then judgment-c doesn't demand any such confidence *either*: again, the *aim* at aptness is one thing and your confidence in that aptness is another matter.

The question here is whether to target the *aim* an agent has in judging when we are trying to define judgment in a way that excludes guessing. But if the difference between judging and guessing really just is a matter of confidence, it doesn't look like the *aim* of judgment is the place we need to focus our attention, in excluding guessing. Both judgment-s and judgment-c might fail to exclude certain forms of guessing, if the difference between judging and guessing really is a difference in confidence. The point here is that it's not the difference *between* judgment-s and judgment-c that will capture the distinction between judging and guessing.

We might still ask if the complex model of judgment is *required* to make better sense of the distinction between guesses and judgments, even if the definition of judgment doesn't get us all the way there by itself. Let's say that we accept that what matters to the difference is a matter of second-order confidence about the aptness of your method of coming to some affirmation on a topic. This level of confidence is obviously more closely connected to judgment-c than to judgment-s. For this reason, we might still have higher hopes for making sense of the distinction between guesses and judgments with the complex model than with the simple model.

Unfortunately, I don't think such hopes are well placed. This would only help us if a difference in second-order confidence could make a difference in whether you're guessing or judging even with no corresponding difference in first-order confidence. But it's very difficult to understand cases where first-order confidence in the actual answer to a question is as strong as it would need to be for judgment while second-order confidence is lacking in this respect. That's partly because, on Sosa's picture, the relevant kind of confidence is playing an operative, occurrent role in what you're doing while you're doing it. If your lack of confidence in your method is so immediately salient to you in guiding what you're doing, why wouldn't it immediately imply a lack of first-order confidence as well? Sosa has indirectly indicated an awareness of this issue in the past, but I'm not sure he has fully appreciated how it can undercut the difference between guessing and judging.⁴⁰

⁴⁰ J&A 75 n14: "Take subjects who feel confident on the first order. They may still feel unsure on the second order, as to how or even whether they have a competence that reliably delivers that first-order assurance. This stance seems possible even if not perfectly coherent." I'm not sure the stance is coherent enough to be possible.

3.3. Summary

This discussion reveals that Sosa's complex model of judgment simply doesn't do the work required to distinguish guesses from judgments. That doesn't yet imply that the complex model is false; it just means that this motivation for the model is not compelling. We shouldn't accept the complex model of judgment for its capacity to distinguish guesses from judgments.

4. Suspension

Sosa has long maintained that suspending judgment on some question must also be assessable by standards of performance. But there's an initial challenge to this approach: "Telic normativity is a normativity of *attempts*, but isn't suspension a paradigm of something that is *not* an attempt?"⁴¹ In this section I'll detail how the way he handles this challenge leads to the complex model of judgment. Then I'll suggest a better way of handling this challenge with the simple model.

4.1. Sosa on the normativity of suspension

Sosa tries to deal with suspension as a performance by identifying a genuine aim that both judgment and suspension share. Within the context of an inquiry into a question, you are acting on an aim which could be met either by judgment or by suspending judgment on the matter. The aim he identifies is a biconditional aim "to affirm alethically (on the given question) if and only if one would so affirm aptly."⁴² To be clear, this does not revise his definition of judgment, which remains the same:

Judgment-c = $_{def}$ an affirmation that *p* made as an attempt to *answer* Qp *correctly and aptly*.

It is rather to say that any instance of acting on this constitutive intention of *judgment* also has to involve, as a "subordinate" aim, this biconditional aim.⁴³ The suspension of judgment on the question at hand can satisfy this biconditional aim, Sosa claims, by avoiding an outcome of inapt judgment when apt judgment isn't clearly available.⁴⁴ For this reason, we can treat the normativity of suspension as the normativity of broadly the same kind of performance as judgment. Then the very same epistemic competence that contributes to judgment-c performed well is that which makes it apt to suspend judgment on some matter.⁴⁵

⁴¹ EE 66.

⁴² EE 55, emphasis original to the text.

⁴³ It's unclear what exactly a "subordinate" aim is, but Sosa does make it clear that it is *not* the relationship between some action type and the means you take to perform that action type. "Action hierarchies are normally ordered through a certain 'by' relation … That is *not* how our epistemic hierarchy need be ordered when we inquire" (*EE* 54-55).

⁴⁴ Strictly speaking, you could meet this biconditional aim by *avoiding* evidence and avoiding any doxastic commitment, but I'll leave aside this point for now. See *EE* Chapters 5-6 for discussion.
⁴⁵ EE 104.

How does this story motivate Sosa's complex model of judgment? The reasons came out more explicitly in *Judgment & Agency*, but they still drive the discussion in *Epistemic Explanations*. I'll summarize the reasoning here.

The way that Sosa ensures that suspension of judgment is treated with *performance* normativity is by casting suspension of judgment as one way to achieve a further aim that you *must* have whenever you are making an attempt at judgment: that further biconditional aim to affirm iff apt. This is a matter of introducing a specific further aim to deal with suspense alongside judgment. Because Sosa deals with suspension by introducing this further aim, we need to make sure that this further aim incorporates the conditions for appropriate suspension. For suspension, the conditions of appropriateness do *not* just have to do with the *correctness* of the judgment you would otherwise make. It's appropriate to suspend just when judgment would not then be *apt*, due to a lack of evidence or something else that diminishes your competence to judge. Thus suspension itself inherently needs a second-order condition for success. But if judgment and suspension must both share an aim—that biconditional aim mentioned above—then judgment, too, must involve some condition concerning aptness in its own success conditions.⁴⁶ This makes the complex model of judgment look inevitable.

This story about performance normativity in suspension is needlessly complicated. There is a much simpler alternative, one that doesn't involve attributing this sophisticated biconditional aim to judgers, and thus one that doesn't make the complex model of judgment look inevitable.

4.2. Instrumental assessment

Here's an alternative way to treat the performance normativity of suspension. First, recognize (i) that judgment involves *means* taken towards its aim, (ii) that any such means can issue in a suspension rather than a judgment, and (iii) the use of a means already makes available a type of performance assessment that of judgment *or* the suspension, namely **instrumental assessment**. In such assessment, we evaluate whether the means taken towards the relevant aim is actually a good means of getting to that aim.⁴⁷ Note that the means is so assessable in the context of acting on that very aim; we don't need to introduce any other associated or 'subordinate' aims to make an instrumental assessment of some attempt.

If you use some means M towards attaining the end of judgment-s—here, the aim of answering the question *correctly*—your whole structured performance is assessable by whether M really is a good means of answering the question correctly, in general. In the context of judgment—considered on the simpler model as judgment-s—M

⁴⁶ See *J&A* 82: "Suspension thus understood involves a second-order intention … Judgment, whether positive or negative, is on a level with suspension, as part of a threefold choice: affirming, denying, suspending. Judgment too thus involves a second-order intention to affirm in the endeavor to affirm aptly … So, an aim that might be shared by epistemic affirmation and suspension is the aim to affirm aptly and only aptly."

⁴⁷ J&A 77ff., EE 68ff. Note that this is primarily a point about constitutive means, although you can run a corresponding point about preliminary or preparatory means. See EE, 22ff.

might be *negatively* assessable because it leads to too many erroneous judgments, or because it leads to too few correct judgments.

Here's an example. Say you decided to figure out whether there's traffic on the bridge by asking your neighbor. Here, your means is asking your neighbor whether there's traffic on the bridge, and this is a means you take to correctly answer the question of whether there's traffic on the bridge. That's a means to judgment-s; here we are considering an alternative model that has no need for judgment-c. We can assess your taking this kind of means by whether or not asking your neighbor really is a reliable method of correctly answering the question about traffic.

Now consider how suspension can arise in using some such means to judgment. Suspension is forbearing from judgment. Sosa recognizes that suspension is necessarily one of the options that faces you whenever you are making an attempt at judgment.⁴⁸ If this is the case, then any means to judgment will also give certain conditions under which you simply don't make that judgment. Then the means is also assessable by these conditions.

To return to our example: your means of answering the traffic question might be a means that leads to suspension just when your neighbor expresses any shyness in her reply to you. If you take her shyness to exclude the possibility of your making a judgment, you might miss some correct judgments—say, if she's needlessly shy about answering your question—and that would make this means to judgment a *worse* means to judgment.

So far, this is just about assessing a means of making a certain kind of judgment, generally speaking. That means is assessable by the standards of judgment, and that means is one that can lead to judgment or to suspension. But how does all this give us a way of assessing an individual instance—of judgment *or* of suspension?

When you take some such means in an attempt at answering some question correctly, what you do is assessable instrumentally. Your judgment is assessable for its success (correctness), its adroitness (the competence so manifested), and its aptness (whether the success is due to that competence's exercise). If the means you have chosen towards trying to answer the question appropriately results in a *suspension* of judgment instead, it is nonetheless still an instance of *trying to answer the question correctly*, and so this very instance of taking this means is assessable instrumentally by how well it leads to successful judgments. It is still assessable in that way because it was an attempt directed ultimately at successful *judgment*, and so the standard by which we can judge this performance intrinsically involves judgment's conditions of appropriateness.

To return to our example one more time: say your neighbor's reply in some particular case is that there *is* traffic on the bridge, but she says this in a shy way, and that leads you to suspend judgment on the traffic question. There's no judgment you make here, but we can still assess your use of the means—which *in fact* led you to suspend—in how it stands as a means to such judgments. All the same considerations about the goodness of the means come into play in a case of suspension.

Note, in particular, that on this alternative picture of the relationship between

⁴⁸ See e.g. EE 60.

suspension and judgment, we can *still* treat the normativity of suspension as the normativity of broadly the same kind of performance as judgment. What's more, it is even more obvious how the very same epistemic competence that contributes to judgment performed well is that which would make it apt to suspend judgment on some matter.⁴⁹

Does this story generate the right result about the appropriateness of suspension? You might at first worry that incorporating suspension in this way would give suspension a *first*-order condition of appropriateness. After all, the means is meant to be a means to *correct alethic affirmation*, on the simple model of judgment we have been considering. So if the means issues in a suspension, isn't that instance of suspension really only assessable by whether it is itself a correct alethic affirmation?

It's not, because it isn't an alethic affirmation at all. Suspension is the actual result of a process which otherwise *could have* resulted in an alethic affirmation, but since it doesn't *actually* result in an alethic affirmation, there is no content to assess for accuracy here. Instead, we must jump up a level to assess the means itself that issued in a token suspension here. This involves ascent to a type-level, modally thick form of assessment that is always available whenever we use means to accomplish our aims. Such instrumental assessment doesn't focus on whether the means accomplished that overall aim in this very instance, but rather focuses on whether the means is overall a good means of accomplishing that aim. Thus the token suspension is assessable in this modally thicker way. The goodness of suspension can be explained in terms of the goodness of the means to successful judgment.

Must a good means of judgment align with appropriate suspension? If the means to judgment is good, the judgments made using this means will be adroit—and, if true because adroit, apt. The means to judgment will be *worse* to the extent that it delivers suspension in cases when the judgment would be adroit. That just would be a condition in which suspension would be inappropriate. We can explain why the suspension is inappropriate directly by pointing out that judgment would have been adroit in these circumstances. That is just to impugn the means taken to judgment. Thus the goodness or badness of the means is directly connected to whether or not judgment would have been *adroit*—and thus could have been *apt*—in context.⁵⁰

Isn't there also a way in which we can *negatively* assess your suspension, when it's the result of a means you took towards actually making a judgment on some matter? It seems worse precisely because it doesn't actually answer the question, which is what you ultimately tried to do.

Yes, we can negatively assess such suspension in this way, but this is also the right theoretical result. Sosa is concerned to clarify that there is a sense in which any suspension *does not meet* your dominant aim of answering the question at hand.⁵¹

⁴⁹ It also seems to me that this simpler description of the relationship between suspension and judgment helps us better capture the failure involved in premature suspension (*EE* 98), as well as the reasons you might be "doomed" to inadequate performance by lacking a certain competence (*EE* 103). It remains to be seen whether this would help us clarify the sense in which the Pyrrhonian skeptic enjoys no advantage, as Sosa says (*EE* 111).

⁵⁰ Thanks to an anonymous referee for pushing me to clarify this point.

⁵¹ EE 106-111.

This is intuitive, as suspending judgment is not a matter of answering the question at hand, but forbearing from answering the question. What's important to see here is that this negative assessment of an instance of suspension is completely compatible with a *positive* assessment of the means that happened to issue, here, in suspension. We don't fault the means to judgment overall if it is in general a good means to the ultimate aim of making a correct judgment. A means can be a good means of making correct judgments not only despite but also *because* it sometimes issues in suspensions instead.

Sosa is also careful to clarify that there are conditions under which *neither* suspension *nor* judgment is appropriate. That, too, can be explained on the view I am proposing. Your suspension can be assessed instrumentally in terms of the means used to judgment. If that's not a good means to judgment, your suspension is also deemed inappropriate by instrumental assessment. But that need not imply that judgment *would have been* apt; the means you're taking is a bad one, so if it issued in judgment, it also wouldn't manifest epistemic competence. The upshot is that you could only have approached the question at hand successfully via some *other* method of coming to a judgment. This might not even be available to you, thus dooming you to inadequacy in your inquiry into the relevant question.⁵²

Let me close by considering a potential response on Sosa's behalf, suggested by an anonymous referee. Sosa might claim that suspension is appropriate and fully creditworthy in a particular case only if the means that led to this suspension was aptly chosen. But if this is true, and suspension's appropriateness requires this additional condition, the alternative model I'm suggesting no longer looks admirably simple in comparison with Sosa's own treatment of the normativity of suspension.

This response should be treated in much the same way as I treated considerations about quality of choice of attempt above (§2.2). First, there isn't always a choice of means; sometimes you have only one means at your disposal, and sometimes your means is so habitual that you slide right into using it. But it does not in general seem to be a condition on our being able to assess suspension at *all* that we be able to identify and assess an event of choosing that means. Even if a good choice of means makes the suspension that results from that means even more creditworthy, it's not clear that we should adjust the fundamental picture of suspension overall to account for this additional credit available. Even if the credit attaches to the occurrence of the suspension-and not to the prior event of choice of means, when there is onethis does not demand that we see the suspension itself as intended "to affirm alethically (on the given question) if and only if one would so affirm aptly."53 The assessment of the choice of the means is one thing, and the aim taken up that results in suspension is another thing, and assessing the former does not require inflecting the relevant aim with second-order attention to the aptness of the relevant aim. Thus this response on Sosa's behalf does not go far enough to privilege his picture of the normativity of suspension over the alternative I have begun to sketch here.

⁵² EE 103.

⁵³ EE 55.

4.3. Summary

We can understand suspension in a simpler way than Sosa suggests, without having to attribute to the agent any complicated second-order aims that suspension could execute. Suspension can be the result of a means taken towards an attempt at judgment; insofar as it is the result of a means so taken, it is instrumentally assessable in a way that depends directly on the aim of *judgment*. Since instrumental assessment necessarily ascends to a modal level, it bears directly on whether or not judgment would have been apt, and we can explain the appropriateness of suspension in terms that concern the aptness of judgment.

In short, judgment and suspension can receive performance-normativity assessments in a way that depends on the *same* aim, but that aim can just be the aim of judgment. We can still explain the second-order appropriateness conditions of suspension without incorporating anything second-order into that aim of judgment. The second-order nature of suspension emerges from the second-order nature of instrumental assessment, rather than some second-order content of an aim towards which both judgment *and* suspension strive.

If all this is correct, then explaining the nature of suspension doesn't require that we accept a complex model of judgment rather than a simple model of judgment.

5. Conclusion

Sosa commits to a complex model of judgment, on which judgment is just this:

Judgment-c $=_{def}$ an affirmation that *p* made as an attempt to answer *Qp* correctly and *aptly*.

He has offered three kinds of arguments why we need this over a simpler model, on which judgment is as follows:

Judgment-s =_{def} an affirmation that *p* made as an attempt to *answer Qp correctly*.

This paper has argued that none of these arguments succeed.

The first of these arguments claimed that we need a complex model of judgment to exclude knowledge-spoiling luck from apt judgment. But the careful amendments to Sosa's view offered in *Epistemic Explanations* undercut this argument, and there isn't any other clear luck-related reason to prefer a complex model of judgment to a simpler one.

The second of these arguments claimed that we need a complex model of judgment to differentiate judgment from mere guessing. But there are guesses that seem to match the profile of judgment-c quite nearly, and it's not clear what the complexity of judgment-c would contribute to ruling out these guesses.

The third of these arguments claimed assessed the performance quality of suspension by introducing a substantive second-order aim to cover both judgment *and* suspension. This would require judgment to have a second-order aim in itself.

But there is a more straightforward way to assess suspension of judgment using performance normativity by understanding instrumental assessment of the means we take in attempting judgment.

I conclude that none of these arguments gives us good reason to prefer Sosa's complex model of judgment over the simpler one I have sketched here. Given that the complex model of judgment has its own costs, we should instead prefer the simpler model of judgment.

Acknowledgements redacted for blind review.

References

Sosa, Ernest (2007). A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I. New York: OUP.

Sosa, Ernest (2009). "Knowing Full Well: The Normativity of Beliefs as Performances." *Phil Studies* 142: 5-15.

Sosa, Ernest (2015). Judgment & Agency. New York: OUP.

Sosa, Ernest (2021). Epistemic Explanations. New York: OUP.