

ARTICLE

Suspension as Spandrel

Ernest Sosa*

Rutgers University

*Email: sosaern@gmail.com

Abstract

A telic virtue epistemology was presupposed in our treatment of insight and understanding. What follows will lay out the main elements of that telic theory and explore how it provides an epistemology of suspension.

Keywords: Virtue epistemology; suspension; telic normativity; knowledge; aptness

1. A virtue epistemology

1.1. All achievements are bound to be attempts that succeed, but the converse is false. Success by luck rather than competence need not be achievement. However, there are many ways in which a successful performance can be lucky without falling short as an achievement. In particular, there are many ways to be lucky in affirming correctly while still achieving epistemic success, thereby knowing the truth of what you affirm. You might be lucky to be alive, for example, having barely escaped a bolt of lightning. You might be lucky to retain the proper use of your faculties, as you drink by luck from the one safe glass out of many before you. Despite all this, you can still know things through the exercise of faculties retained by luck.

For virtue epistemology, the luck that blocks knowledge on a given question is luck that clashes with the subject's getting it right through competence on that question,¹ a clash that is not distinctive of the epistemic domain. *Achievement*, whether epistemic or of any other sort, is blocked by such luck. An attempt constitutes an achievement only if it succeeds sufficiently through competence.²

1.2. Achievement requires success that is *apt*: through competence rather than luck. This emerges from a review of telic theory's five main phenomena: *attempt*, *success*, *competence*, *aptness*, *achievement*.

Archery provides a vivid example of that structure. If an archer shoots aiming to hit a certain target, we can assess that shot in various respects. First, does it *succeed*? Does it hit the target? Second, how competent is that shot? The arrow may exit the bow with an orientation and speed that would normally take it straight to the bullseye. Even if a gust diverts it, that shot might still be highly competent. It can be *adroit* without being *accurate*. And it can be accurate by luck, without being adroit. But even a shot that is both accurate and adroit might still underperform. An arrow adroitly released from a bow may be headed straight to the bullseye when a gust diverts it, so that it would now miss the target, except that a second gust puts it back on course. The archer

¹This is a first approximation.

²Thus the apt title of John Greco's *Achieving Knowledge* (Greco 2010).

succeeds in that attempt to hit the target, and the shot is also competent, as the arrow leaves the bow perfectly directed and with the right speed. But the shot is accurate because of the lucky second gust, not by manifesting its adroitness. So, it is accurate and adroit without being *apt*.

When generalized to all attempts, of whatever sort, that is an account of the telic normativity of attempts as attempts, in terms of their accuracy, adroitness, and aptness.

1.3. *Telic* credit is earned when a success is thus “creditable” to the agent, *attributable* to them, without necessarily importing any more substantive axiological standing. A shot that constitutes a “perfect” murder might be an excellent shot, one thus creditable to its agent, while constituting an abominable crime, without earning any *moral* credit.

1.4. If a driver is both skillful and in good shape, does that make him competent to drive safely on a given road? Yes, if the road is in appropriate condition. But what if the road is then entirely covered with oil? That surely deprives the driver of the *complete* competence to drive safely on that road. In acknowledging this, we recognize varieties of competence. The complete competence requires not only that the agent be sufficiently skilled and in proper shape, but also that the situation be favorable as well.

Similarly, might an *archer* be skillful and in fine shape while still lacking the complete competence required for success through competence, because spoiler wind is too likely? Suppose that, although high wind is *very* likely, in fact no gust intervenes, and the arrow goes straight to the bullseye, just as it would normally do, given its orientation and speed off the bow. How does the high likelihood of spoiler gusts bear on the quality of that shot? Does it put the archer in an *inappropriate* situation? Does he thereby lack the competence required for *apt* success?

Surely the archer deserves a high measure of credit so long as his arrow is *in fact* unaffected by wind on the way to the target, no matter how likely a spoiler gust may have been.

That being so, the “situation” required for that high measure of credit is clearly not a *modal* property of the spatiotemporal volume involved. Success in hitting the target across the relevant space is quite unlikely at that time, despite our archer’s excellent skill and shape. What makes success so unlikely is the high risk (by hypothesis) of a spoiler gust. However, so long as no spoiler gust *in fact* comes along, our archer enjoys the competence required for highly creditable, *apt* success.

Surprisingly, aptness depends not at all on a modally safe situation. It depends rather on an *actually* favorable situation, no matter how accidentally it may be favorable. Beyond this, the Situation need only combine with the agent’s Skill and Shape to make *that particular* SSS combination likely enough to yield success.³

Apt performance, including apt epistemic performance, is not dependent on how safely one *possesses* relevant competence. This applies to all three sorts of competence: first, the (innermost) skill; second, the skill plus the required inner shape; third, the skill and shape, in turn, plus the required situation. None of these varieties of competence need be *safely* in place. The safety that *does* seem required for apt performance, including apt judgment and belief, is rather the SSS-relative safety constituted by the fact that one is (actually, however luckily) SSS-competent enough, so that, if one tried *when thus* SSS-competent, then one would likely enough succeed.⁴

Whether lucky or not, *possession* of each pertinent S facilitates corresponding competence. However, in other ways luck does clash with competence. How much a success is by *credit-reducing* luck depends on how little it is due to competence. *Excess of such*

³Still in first approximation.

⁴To continue with our first approximation.

luck aligns with *deficiency of corresponding competence*, which reduces or blocks relevant credit to the agent.

1.5. Here again are the five main ideas of our account: they are those of *attempt*, *success*, *competence*, *aptness*, and *achievement*.

Archery not only provides an example of a telic triple-A normativity constituted by those five ideas. It also shows how achievement comes in degrees within at least two dimensions. One dimension is that of the apt shot, accurate *because* adroit. The other is that of the *fully* apt shot, where the agent aims not just at accuracy but at aptness and succeeds through competence in this more complex endeavor.

As Diana surveys a landscape in search of game, she may see prey in the distance (in good light and calm wind). If the prey is beyond her range, then her shot would be poorly selected. But if the prey is within her range, then her shot might be well selected, one she'd be well advised to take, given her aim to hunt well. If a shot is too risky, it is ill-advised. A shot can attain quality *in the specific regard of being well selected*. A well-selected shot can thus rate higher in that regard than one that falls short through pertinent negligence or recklessness.

A dimension of second-order evaluation of Diana's shot thus involves more than its aptness, its success through (first-order) competence. Also relevant is whether the attempt is well selected so as to avoid recklessness, and even negligence.

That is so even if the dimension of interest to us is that of creditability, independent of moral concerns such as whether the act is a murder, and independent of the admirability that involves degree of difficulty. An instance of absolutely certain knowledge, of the highest degree of relevant epistemic quality, need involve no difficulty *at all*. Think of the *cogito!* Not much of an "achievement," that one, but still a minimal one nonetheless. Or, rather, not much of an achievement *in respect of the minimal difficulty overcome*. But still an achievement that reaches epistemic heights nonetheless. How so?

When successfully enough guided that way, an attempt rises to the level of the *fully apt*. Nothing short of this will suffice for *achievement full well*. 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. And this sort of luck reduces or blocks credit to the agent for their success, as it reduces or blocks credit to the agent for the aptness of their success.

1.6. Going beyond virtue theory in general, here is a main thesis of virtue *epistemology*:

that the normativity of *knowledge* is a special case of such *telic normativity*.

Knowledge is thus a central sort of epistemic achievement. Here we find the traditional issues of skepticism, and other issues of the nature, scope, and value of knowledge.

Gettier protagonists now seem to fall short either because their pertinent belief falls short of aptness altogether, or because it falls short of *full* aptness.

1.7. In *functional* epistemic performance, including functional belief, the agent aims at truth and aptness of representation, but does so only implicitly and teleologically, as when our perceptual systems aim at correctly representing our surroundings.

Here, however, we focus mainly not on functional, subconscious representation, but on judgment, and on judgmental belief, where the agent aims *with conscious intention* to get it aptly right on a given question. What is involved in this?

Sayings include utterances of declarative sentences by actors on a stage. When an actress says "I am the queen," she does not really say *that* she is the queen. If Queen

Elizabeth is in the audience, she would be wrong to object: “No, she isn’t!” Play-acting is pretense, and pretend affirmations are not the genuine article. Only some sayings of declarative sentences are *affirmations*, genuine sayings of the form *that p*. And only some affirmations are *alethic* affirmations, in the endeavor to get it right on a given question (whatever else one may also be endeavoring to bring about thereby).⁵ An alethic affirmation might be just a guess, as when a contestant tries to affirm the correct answer to a quiz show question. But an oncologist would aim not just to guess but to affirm competently, indeed *aptly*. Only an alethic affirmation can amount to a judgment, which it can do only if it aims not just at truth but also at aptness. This yields the following hierarchy.

Saying: of “p”

Affirmation: saying that p

Alethic affirmation: endeavor (attempt) to get it right by affirming that p

Judgment: endeavor (attempt) to get it right *aptly* by alethically affirming that p⁶

These can all be public, in outer speech, or private, in silent soliloquy. They are commonly and generally free acts.

2. The place and importance of suspension

2.1. In the domain of action in general, not just epistemic action, we find a “forbearance” that amounts to *intentional omission*. Here two varieties can be distinguished through the following formulation:

Forbearing from X’ing in the endeavor to attain an aim A.

Where might parentheses go into that formulation? Here are two options:

Narrow-scope: (Forbearing from X’ing) in the endeavor to attain a given aim A.

Broad-scope: Forbearing from (X’ing in the endeavor to attain a given aim A).

⁵True, there is something that the actress playing the queen is doing right when she says “I am the queen,” which the actor playing the jester would not get right if he voiced the same sentence. The “jester” would pretend incorrectly, would not speak his lines correctly. The “queen” would at least pretend correctly, would speak her lines as expected. But her pretend assertion that she is the queen would not be a real assertion, one that might call for correction by Queen Elizabeth as she sits in the audience.

⁶*Objection:* It’s not clear that judgment can be defined in terms of the endeavor to get it right aptly by alethically affirming that p. After all, conjecturing is distinct from judging and yet when conjecturing that p one will aim to get it right aptly by alethically affirming that p. *Reply:* Yes, so far our picture is a first approximation. In a fuller development, we distinguish varieties of judgment. There are of course varieties of *expert* judgment more ambitious than ordinary judgment. What is more, even an “educated guess” might count as a sort of judgment, since one does not then affirm without attempting to affirm with alethic aptness. The educated guess does attempt to affirm with some degree of aptness, constituted by some degree of competence. And similarly for conjectures that go beyond sheer guessing. Bottom line: judgment that p can be defined as endeavor to get it right aptly (aptly outright or aptly full stop) by alethically affirming that p; *but* we must recognize the varieties of judgment determined by the degree and the sort of aptness aimed for. We encountered in Chapter One the idea of *firsthand* judgment, and we’ve also encountered the distinction between the minimal or zero aptness aimed for by the quiz show contestant and the degrees aimed for by the expert oncologist.

2.2. That distinction among forbearances enables us to zoom in on the suspension of judgment that is of main epistemic interest. But first consider how suspending comes in at least two varieties.

Epistemically idle suspending is intentional forbearing from judging (from both positive and negative judging) on a certain question $\langle p? \rangle$, *without* doing this in pursuit of any ulterior epistemic objective concerning that question. By contrast, *epistemically aimed* suspending is subordinate to a particular kind of intellectual attempt: namely, the attempt to affirm alethically on the given question (positively or negatively) if and only if one's alethic affirmation would be competent and indeed *apt*.

The suspending of main interest to us is a sort of forbearing, but it must be specifically *narrow-scope* (narrow-scope intentional double omission of alethic affirmation, both positive and negative). And it must be epistemically aimed suspending. This is because the theory of knowledge *presupposes* a focus on *inquiry*, in the *broadest* sense, where the subject simply takes up a question, which you might do just casually and implicitly, as when you walk down the street while constantly monitoring your surroundings.

Judgmental knowledge involves judgment. You aim to get it right aptly on a given "whether" question $\langle p? \rangle$ by affirming alethically on that question (positively or negatively). To succeed (aply) in this aim is to judge with (apt) success.

When we "consider" such a question, we may do so implicitly, as with our implicit monitoring in our walk down the street. Alternatively, we may focus consciously on our question, and ponder how to answer it, if at all. Either way, we may be led either to judge (positively or negatively) or else to suspend judgment, and these (judging and suspending) then share an aim: *to affirm alethically if and only if one would affirm aptly*. Judgment, recall, aims not just at the truth of one's affirmation, but at its aptness. When one takes up a question, when one inquires seriously, as opposed to just guessing, one aims for alethic affirmation that will be apt, not *just* true (not just true and perhaps *inapt*).⁷ One aims like the oncologist, not like the quiz show contestant.

When one faces *judgmentally* a question whether p , one deliberates on whether to affirm alethically (positively or negatively) or suspend (intentionally omitting alethic affirmation). Judgment on whether p would require aiming for *apt* alethic affirmation. So, competent pursuit of that aim would require aiming to affirm only if one (likely enough) would affirm aptly. One puts oneself in the appropriate shape and situation and approaches the question with the required skill so that one affirms only if one would do so aptly. Aiming to satisfy that conditional is inherent to a sort of serious inquiry defined in part by so aiming. And this is half of our biconditional objective: to affirm (positively or negatively), on the question whether p , if and only if one would thereby affirm aptly. (In abbreviated form, the objective is: to affirm re $\langle p? \rangle$ iff one would affirm aptly.)

It is crucial in epistemology to distinguish theory of knowledge from theory of inquiry. Whether to take up a question, which is a concern absolutely central to the normativity and theory of inquiry is quite irrelevant to the normativity and theory of whether one knows through alethic affirmation on that question. One can know with great certainty and utmost epistemic standing truths whose utter triviality makes it quite inadvisable to inquire into them. The lowest justification of inquiry into a question leaves it entirely possible that one attain the very highest level of epistemic certainty for one's knowledge of the right answer to that question. This is why it is so important to distinguish between narrow-scope suspending and broad-scope suspending.

⁷And the like holds as well when one just engages a question implicitly, without consciously intentional pondering.

Why think that in such inquiry one must adopt the subsidiary aim to affirm re $\langle p \rangle$ *only if* by so affirming one would affirm aptly? If one affirmed when by so affirming one would affirm *inaptly*, then one would err in a way highlighted by Descartes, as a deeper error than mere falsity. One of course errs when one's alethic affirmation is false. But one is also in Cartesian error when one's alethic affirmation is true but inapt. Such affirmation is true by a sort of luck that blocks credit to the thinker. The attainment of truth is then not through competence but through that sort of epistemic luck. In proper inquiry we aim not just to attain apt alethic affirmation on the question addressed, whether p . We aim also for the avoidance of inaptness, of Cartesian error.⁸ True, if we attain aptness, we thereby avoid inaptness. But one might avoid inaptness without attaining aptness. So, even when one *fails to attain* one's dominant aim of inquiry, one might *still attain* a subsidiary aim, namely avoiding inaptness, which earns partial credit of a sort.

Unless aimed at least in part at avoiding inaptness, deliberation 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 properly guides oneself to do so, one's attainment of aptness is relevantly lucky. It is *insufficiently* owed to the exercise of reliable enough pertinent competence.

So much for half of our subsidiary aim involved in the sort of serious inquiry highlighted here, according to which one is to affirm *only if* by so affirming one would affirm aptly. What of the other half of our biconditional aim? Why think that in inquiry one must adopt the subsidiary aim to affirm re $\langle p \rangle$ *if* by so affirming one would affirm aptly? Well, isn't it just inherent to such inquiry that one should affirm once one has put oneself in a condition wherein if one affirmed one would do so aptly? Isn't this just part of what it is to aim for aptness of affirmation in one's attempt to answer the question into which one inquires?⁹

⁸*Objection:* This does not generalize to performances in general. When playing a game of darts in which the first to hit the bullseye 20 times wins, I might be best advised to take many shots quickly, without regard for avoiding failure or inaptness. *Reply:* This objection prompts recognition that attempts can be nested. One attempts to hit the bullseye 20 times as quickly as possible, and in pursuit of that dominant objective one may make a series of quick attempts without trying for high competence or aptness with each attempt, but while trying to attain the *quick* 20-fold success with the highest available competence and aptness. So, the subsidiary attempt to hit the bullseye by shooting quickly properly reduces the degree of competence and aptness aimed for in that *particular* shot, for the sake of producing a *sequence* of shots that will thereby attain *its* success more competently and aptly. And this reveals a way in which an encompassing dominant attempt can trump a component subsidiary attempt, so as to lower the degree of aptness and competence properly aimed at in the latter. (Further afield, an extension of this point will also apply to the proper subordination of individual attempts in deference to important enough collective attempts, since humans are not only rational but also properly *political* animals.)

⁹Here is the relevant objective when I face the question whether p :

- O I affirm (positively or negatively) on that question if and only if I thereby affirm aptly.
- O is logically equivalent to the following:

O+ & O (I affirm positively on that question if and only if I thereby affirm aptly) & (I affirm negatively on that question if and only if I thereby affirm aptly).

So, O is logically equivalent to (O+ & O-).

Now, to affirm positively on the question whether p is to affirm that p . And to affirm negatively on the question whether p is to affirm that not- p . So, (O+ & O-) is equivalent to the conjunction of the following two propositions:

2.3. A desirable level of human knowledge is the *fully apt alethic affirmation*. Such knowledge constitutes a desirable sort of success in inquiry.¹⁰ It thus provides a (main) norm of judgment, whether public or private. And it is thus not only a norm of *judgment*, but also a norm of *suspension*. The immediate aim that one promotes when one suspends properly is the aim of affirming alethically if and only if so affirming would be apt. This subsidiary aim is one that must be attained aptly if one is to attain aptly one's dominant aim. The subsidiary aim of proper suspension is that of affirming alethically if and only if one would thus affirm aptly. That objective is one properly pursued in suspending. And it is an aim subsidiary to hierarchically superior aims: that of apt alethic affirmation (animal knowledge), and that of apt judgment (knowledge full well).¹¹

2.4. It may be argued that the proposal here overintellectualizes the normativity of ordinary belief. However, the biconditional aim proposed as inherent to proper belief and proper suspension need not be a conscious aim.

Consider the *omissions* that might be made in functional, teleological pursuit of such an aim. Such omissions too can remain on a functional, teleological level.

Take for example the plover hen who seeks to protect her nest through her broken wing ruse. She thus seems to aim at issuing her ruse iff a predator gets too close. In cases where a predator is too far, her omission of the ruse need not be owed to inattention, temporary blindness, or any other negligence or flaw. Rather, she may properly omit the ruse in pursuit of that biconditional aim. What is more, sometimes the plover hen can tell whether the nearing animal is likely to pose danger and she will then use the ruse if and only if she perceives real danger. If this sort of intellectual performance is not functionally beyond plovers, it is not plausibly beyond humans.

2.5. *How is that epistemic hierarchy ordered?* Action hierarchies are normally ordered through a certain "by" relation. You may for example aim to alert a co-conspirator

O+ I affirm that p if and only if I thereby affirm aptly.

O− I affirm that not-p if and only if I thereby affirm aptly.

Now, 'affirm' here has been short for 'affirm alethically' or 'affirm in the endeavor to affirm correctly, with truth.' So, the two conjuncts here are equivalent respectively to:

O+ I affirm-alethically that p if and only if I thereby APTLY affirm-alethically that p.

O− I affirm-alethically that not-p if and only if I thereby APTLY affirm-alethically that not-p.

And each of these is of the form:

A I \emptyset if and only if I thereby \emptyset aptly.

And here \emptyset ranges over first-order attempts, where you act in the endeavor to attain an aim, as when you affirm alethically, where affirming alethically is affirming in the endeavor to affirm correctly, with truth. It is only such attempts that can conceivably fall under the AAA/SSS structure. Only such attempts can be accurate (successful), adroit (competent), and apt (where the success manifests the pertinent competence exercised by the agent).

So, properly suspending requires the intentional adoption of two objectives of form A above, namely O+ and O−. And it requires further that one intentionally OMIT affirming-alethically that p and also OMIT affirming-alethically that not-p, and that one do so BECAUSE of one's intentional adoption of O+ and O−.

¹⁰And this is so whether the inquiry takes the form of conscious pondering or that of implicit functionally teleological "processing."

¹¹More generally, achievement is the norm of attempt. A Hail Mary attempt is one where the risk is judged (correctly) to be extremely high. The guess of the quiz show contestant is an example, as is the shot seconds before the buzzer from way back nowhere near the basket. The first-order norm of an attempt is its aptness. But a higher-order norm is met when the attempt succeeds *full well*, where its aptness is *also* aimed at, and attained aptly. (And so on, up the orders of aptness at least until the agent's capacity runs out. And even before we reach our limits of capacity, the returns may diminish rapidly.)

by illuminating a certain window, aim to do this in turn by illuminating the room with that window, and aim to do *this*, in its turn, by flipping a certain light switch. How do you implement that plan? As follows: By flipping the switch, you thereby illuminate the room; by illuminating the room, you thereby illuminate the window; and by illuminating the window, you thereby alert the co-conspirator. And that is the hierarchical plan whereby you alert the co-conspirator.

That is *not* how our epistemic hierarchy can be ordered when we inquire. Suppose one has to *suspend* as one aims for this biconditional: *to affirm alethically (on the given question) if and only if one would so affirm aptly*. Suspending here is intentionally omitting alethic affirmation (whether positive or negative), in the endeavor to satisfy the two conditionals. If this is subordinate to knowing the answer to that question, it cannot be in the way familiar from ordinary action hierarchies. It can't be that by thus suspending one would *thereby* know the answer, either on the first order (with apt alethic affirmation, or "animal" knowledge) or on the second order (with apt judgment, or "reflective" knowledge full well). Such suspending can never constitute either apt affirmation or apt judgment. Instead it can be a sort of spandrel. It can be an inevitable byproduct of the serious pursuit of a dominant aim. That is what happens when suspending is the outcome of the inevitable pursuit of a subsidiary aim, one that must be pursued given one's primary aim. The subordinate epistemic aim must be *pursued*, for the sake of attaining a dominant aim, namely that of apt alethic affirmation.

Here next is an example that illuminates our unusual way in which an aim can be subordinate to a hierarchically dominant aim.

Suppose I play chess regularly with my young grandson, and I can see the fast improvement, but also the low confidence and the nervous anxiety. I continue to play so that (as I fully expect) he may start to win more and more of our matches and gain confidence and enjoyment in the game. Suppose that to be my master intention as we play our next game. I aim to lose by playing, confident that he is now the better player, so that by playing him I will likely further his winning. But I want the game to be real and fair, so I play to win, but with the overarching and guiding intention to lose. If I did not aim to lose I would not play, since I've been fearing that I might damage his fragile confidence and discourage him to the point of quitting the sport altogether.

I do try to win. And I might of course attain this aim aptly, through the exercise of competence. But I also aim to lose, and we can suppose that, if I attain *that* objective, I will do so competently. It will not be a fake loss. That is not what I want. Rather, I mean the loss to be a real loss. Although I really try to win, yet my *overarching* aim is to *really* lose, which requires that I seriously try to win. Paradoxically, it appears that one can appropriately and seriously aim *both* to win *and* to lose. How can this be?

How? In the way of our example, with one aim subordinate to the other. But this subordination is unusual. A more usual way in which one aim is subordinate to another is through one's *intention* to attain the former, by attaining which one will, *thereby*, attain the latter. In our unusual case, one aim is subordinate to another through one's aiming to attain the latter in part by *hosting* the former. Thus, it is not that I aim to lose to my grandson by winning. Rather, I aim to *really* lose to my grandson in part *by seriously aiming to win*.¹²

¹²*Objection*: "Why not say instead that your aim is to win, though you hope to lose, or would prefer to lose? It is hard to see how one could aim both to win and to lose." *Reply*: What would be incoherent is to

If it happens that I do win, this will be a sort of spandrel. In the circumstances it may be an inevitable consequence of my pursuit of an aim whose serious pursuit was required for the attainment of my dominant aim. And this is the sort of spandrel that suspension turns out to be when it occurs in the pursuit of an aim that must be pursued if one is to attain a certain dominant aim. Here one aims to affirm if and only if affirming would be apt (and otherwise to suspend). This subsidiary aim must be adopted when one pursues a dominant aim to affirm aptly on the question addressed.

In that way, knowledge is a main norm (a hierarchically dominant aim) not only of judgment (and assertion) but also of suspension.¹³

3. Appendix: a broader view of suspension

3.1. We have focused in the foregoing on a particular sort of suspending: namely, occurrent judgmental suspending.

Suspending comes in several varieties, however, determined by two divides. I mean first the divide between occurrent and dispositional suspending, and second the divide between judgmental and functional suspending.

To suspend judgmentally is to *consciously* and intentionally forbear from judgment (either positive or negative) on the question considered, whether p. Functional suspension, by contrast, is constituted rather by the agent's intentionally though *subconsciously* forbearing from functional (aimed) belief and disbelief.

The discussion above has focused mainly on just one of those forms of suspension. We have discussed the consciously intentional omission of judgment, whether positive

aim to win and by *winning thereby lose*. That is too absurd to countenance as an aim that a rational agent could possibly have. But that was not my situation in playing my grandson. I was aiming to lose. I was not just hoping to lose. I was actually trying to *bring about* my grandson's proper and serious win. And of course I realized that such a win would come about only through my seriously playing to win. But still it just seems right that I am playing seriously to win and doing so because I realize that only with such serious play would the win of my grandson be real, not fake. If I had thought that my serious play to win would actually succeed surely enough, I would *not* have played. I would have told my grandson that we should postpone our game and that he should play his playmate instead.

¹³In recent years Jane Friedman has written a series of papers in which she develops a novel account of suspension, in relation to questions, beliefs, and inquiry. (In "Why Suspend Judging" (Friedman 2017), for example, and "Inquiry and Belief" (Friedman 2019).) Although not incompatible, our views are radically different. The main differences derive from the telic character of my own approach. We have seen how this opens up a view of epistemology as an account of one more domain of human endeavor. In this domain of epistemology there are several distinctive aims, so that an array of interrelated attempts unfolds, displaying crosscutting varieties. Some attempts are consciously intentional, while others are functionally teleological. For a second distinction, animal beliefs (mere alethic affirmations) fall hierarchically under reflective judgments, whose aim is not just truth but aptness, or truth through competence. Friedman relates suspension (as she views the phenomenon) to questions, to beliefs, and to inquiry. One important difference between our approaches derives from how telic threefold multiplicity falls into three categories: in how we can approach a question, in the sorts of objectives coordinate with any such approach, and in corresponding forms of inquiry. To take just one distinction, we may have a first-order conscious aim to answer a certain question correctly, as might a quiz show contestant; alternatively, we may aim to answer our question *aptly* (and not just correctly), as might an oncologist. But, again, that is just a partial account. In fact, there are, I believe, various constitutively different sorts of inquiry and suspension, corresponding to the varieties of epistemic aims. Each such aim will import various distinctive ways in which one can forbear from affirming. And suspension in all its varieties I take to be fundamentally a kind of intentional forbearing (or, more generally, a kind of *aimed* forbearing, including two quite distinct approaches to questions: first, involving consciously intentional aiming, and, second, involving just functionally teleological aiming).

or negative, on a question taken up. What follows supplements the foregoing by discussing suspension more generally, to include the functional and not just the judgmental, both in occurrent and in dispositional forms, and both the instantaneous and the deliberately stable and longstanding.

On the first order, suspending is comprised of *absences*: the absence of affirming/believing and the absence of denying/disbelieving. And this is compatible with the presence on the first order of a certain level of confidence, not necessarily conscious, one below some threshold above 0.5 and above some threshold below 0.5. That the subject hosts that level of confidence is not determinative of whether he has the disposition to affirm, however, nor of whether he has the disposition to deny. And it is *these* dispositions that determine whether he relevantly “believes” or “disbelieves,” on our conception of these states: i.e., whether he assents or dissents, or whether he would affirm or would deny. So, it remains that on the first order there is not much, if anything at all, that constitutes the suspending. Of main relevance on the first order are simply *absences* – the absence of believing and the absence of disbelieving – that is to say, the absence of a relevant disposition to affirm and the absence of a relevant disposition to deny.

However, is there not *ever* anything that positively constitutes suspending? Might it not be simply a credence within the thresholds of belief and disbelief? This would be a credence that lies between the disbelief threshold (a threshold below 0.5) and the belief threshold (a threshold above 0.5). Intentions are still involved here, however, whether they are *acts* of decision or choice, or *sustained* intentions that may involve such acts (either by deriving from some such acts, or by being intentions – or, better, policies – to perform such acts in certain conditions; or both). It is one thing to host a credence of a certain degree (with that degree of confidence), and quite another to host a credence of such intensity *while that degree lies within the relevant interval, above the disbelief threshold and below the belief threshold*. This latter requires the enclosure of that degree of confidence within those thresholds. But its being so enclosed depends on the locations of the thresholds.

And what determines those locations? Well, suppose first that the belief threshold is *defined* as the threshold at or above which one affirms or is disposed to affirm if one then endeavors to affirm aptly. Suppose, second, that the belief is *judgmental*, and that the disposition to so judge takes the form of a *policy*. If this policy is sustained freely, as policies are generally sustained, then the being enclosed of that degree of confidence within those thresholds is determined freely, through the free determination of where one begins to affirm and of where one stops denying. But if one freely determines where one starts affirming (or becomes willing to do so, or by policy would now do so), then plausibly one freely determines when one does *not quite yet* affirm (or become willing to do so, or by policy would not yet do so). And if the free determination of where one starts affirming implements an intention to do so, then plausibly the free determination of where one does *not yet* affirm *also* implements an intention, the intention *not yet* to do so. When one *deliberates* on an option and one has not yet decided, this lack of decision is no doubt intentional, and plausibly implements an ongoing intention not yet to decide, and not to do so until the balance of reasons adequately supports one’s decision.

Objection: “When one deliberates without yet having opted among affirming, denying, or *suspending*, one *already* intentionally omits affirming and omits denying. So, this intentional double-omission cannot be what amounts to suspending.”

Reply: Yes, good point, which brings out that there are varieties of suspension. We here have reason to distinguish between *provisional suspension*, which occurs when one is still deliberating, and *conclusive suspension*, the settled attitude that derives from one’s concluding deliberative inquiry by opting for settled suspension. These are thus two

different forms of suspension, both of which we should recognize. One might of course end inquiry without opting for any such settled attitude. In this case one *defers*. One might even defer indefinitely, as one reaches a point where it seems useless to continue inquiry, since one sees no way to settle into one of the three attitudes. One ends inquiry without being prepared to *conclusively* affirm, deny, or suspend. One reaches *no settled attitude* whatsoever, and simply redirects one's attention.

3.2. Settled, conclusive suspension is hence a kind of intentional double-omission. What constitutes such an intentional omission? Say one faces an option whether to wiggle a certain finger. One might then decide to do so, whereby an intention is born, one soon implemented. One might of course decide *not* to do so. However, might one simply intentionally *not* wiggle that finger, with no benefit of corresponding intention? Not plausibly: after all, what makes an omission intentional if not the intention to omit? Of course, an intention need not be at the focus of one's consciousness in order to so much as exist. An intention can be subconscious, as when one tries to put on the brakes in response to the sight of a red light. Even if the brakes are stuck, one tries to put them on, and it is hard to see what might give positive substance to the trying other than the relevant present-directed intention, even if it is not a conscious intention. It might be denied that in such a case one does really try, but how plausibly?

The foregoing suggests that omissions (and actions more generally) can be intentional in either of two ways at least. One way is for the omission to correspond to a present-directed intention, whether conscious or subconscious, an intention to persist in omitting. A second way is for the omission to derive from a past decision to omit indefinitely, even if the continuing omission is no longer accompanied by a present-directed intention to keep on omitting. The continuing omission cannot remain intentional, however, absent any continuing intention to omit, *unless* it continues to be mnemonically owed to the earlier decision to omit. Suppose the relevant mnemonic efficacy lapses, so that now the omission continues only because the whole matter is forgotten, and one simply fails to concern oneself with the relevant options at all. At this point one simply omits, without this being at all intentional. And at this point one *also* stops suspending.

Epistemically most interesting, I suppose, is the suspending that involves a double-omission made intentional through obedience to a continuing second-order intention to omit. But the present approach need not rigidly insist on that. It can allow the intentional omission to be owed to a past rather than a present second-order intention. In order to preserve the significance of the reflective order one need only argue that often enough the intentional character of the double-omission does properly derive from a present-directed intention. This will suffice, especially if in its absence the intentional character of the double-omission must *still* derive from a reflective-order intention. Even when such a second-order intention acts at a temporal distance, it must still have been present in the relevant epistemic etiology, so that the present suspending can derive its epistemic standing from this reflection-at-a-temporal-distance.¹⁴

References

- Friedman J. (2017). 'Why Suspend Judging?' *Noûs* 51, 302–26.
 Friedman J. (2019). 'Inquiry and Belief.' *Noûs* 53, 296–315.

¹⁴This paper is part of a conference on my work in epistemology held in Tenerife in the summer of 2018. I am grateful beyond words to all the participants in that wonderful event, and especially to the editor of *Episteme*, Jennifer Lackey.

Greco J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.

Ernest Sosa is Board of Governors Professor at Rutgers University. He has published papers and books in epistemology, including most recently *Judgment and Agency* (Oxford University Press, 2015) and *Epistemology* (Princeton University Press, 2017).