Episteme (2019), **16**, 413–437
doi:10.1017/epi.2019.34

episteme

ARTICLE

When Evidence Isn't Enough: Suspension, Evidentialism, and Knowledge-first Virtue Epistemology

Lisa Miracchi*

University of Pennsylvania, Philadelphia, PA, USA
*Email: miracchi@gmail.com

Abstract

I motivate and develop a novel account of the epistemic assessability of suspension as a development of my knowledge-first, virtue-epistemological research program. First, I extend an argument of Ernest Sosa's for the claim that evidentialism cannot adequately account for the epistemic assessability of suspension. This includes a kind of knowledge-first evidentialism of the sort advocated by Timothy Williamson. I agree with Sosa that the reasons why evidentialism fails motivate a virtue-epistemological approach, but argue that my knowledge-first account is preferable to his view. According to my account, rational belief is belief that manifests proper practical respect for what it takes to know. Beliefs are the only primary bearers of epistemic evaluation since they are the only candidates for knowledge. However, suspension can manifest a derivative kind of practical respect for what it takes to know. Thus, we can explain why the same sort of assessment is applicable to both belief and suspension (epistemic rationality), and why belief has a privileged claim to these properties. Lastly, I'll look at Sosa's and Williamson's treatments of Pyrrhonian skepticism, which treats a certain kind of suspension as the epistemically superior practice, and argue that my account provides a better anti-skeptical response than either of their approaches.

Keywords: evidentialism; knowledge-first epistemology; virtue epistemology; skepticism; suspension

Introduction

Epistemologists spend most of our time discussing the epistemic properties of belief, and rightly so. However, an exclusive focus on belief may obscure or distort the normative landscape. We do well to periodically expand our purview to other states and events that deserve epistemic assessment. In what follows I motivate and develop a novel account of the epistemic normativity of suspension of belief. Not only is suspension regularly the object of epistemic assessment – as rational, irrational, justified or unjustified – but throughout history it also has often held an epistemically privileged position as the purportedly appropriate response to skeptical concerns. Investigating the normative properties of suspension will also illuminate those of belief, for (typically) if it is inappropriate to form a belief on a question then it is mandatory to suspend, and vice versa. Delineating the contours of rational suspension will thus require us to reexamine those for rational belief.

As will shortly become clear, this paper is heavily influenced by Ernest Sosa's recent work on suspension, especially his paper "Suspension and Evidentialism" (Sosa

Forthcoming), in which he argues that investigating the epistemic properties of suspension motivates the rejection of evidentialism and the adoption of a virtue epistemological approach. I agree in broad contours with both his negative and positive claims, although I develop them in different ways and ultimately arrive at a different kind of virtue epistemology of suspension.

In section 1, I lay out a variety of options one might take in aiming to illuminate the rational assessability of suspension. One possible option, central to Pyrrhonian and other skeptical approaches, is that suspension has the epistemic upper hand: either as the default view, or as epistemically safest – the least risky – of the available options. This should be resisted. It is often impermissible to suspend judgment, and an adequate theory of suspension should explain why. The import of this issue is not merely academic: an adequate theory of epistemic rationality should explain why it is mandatory to believe in human-made climate change or systemic racism, for example. Another view, accepted by Ernest Sosa (Forthcoming), is that suspension is rational *in the very same way* that beliefs are rational (*Rational Parity*). I briefly note a concern about this approach (deferring substantive discussion to section 3) by way of introducing the third option that I endorse, on which suspension has epistemic properties but only derivatively (*Rational Derivativeness*).

In section 2, I investigate the two challenges Sosa (Forthcoming) poses for evidentialism and argue that the first one fails because it depends on *Rational Parity*. I argue, however, that the second challenge succeeds and reveals something deeply right about our epistemic practices that evidentialism cannot account for, but to which virtue epistemology is well-suited. We assess each other, and each other's belief and suspension, not just with respect to what evidence one has, but with respect to what evidence one *should* have. The evidentialist cannot properly account for this legitimate practice.

I further develop this objection, focusing on a particular role that our epistemic assessments should play: they should help us to hold each other accountable for what we believe and why. An assessment of belief or suspension as rational is a sanctioning of the agent's epistemic comportment. It is important that our assessments of an agent's epistemic comportment extend beyond how the agent deals with her evidence to (at least) how she gathers evidence and when she treats a body of evidence to be enough to settle a question. As such, our assessments of epistemic rationality are determined by features that go beyond what evidentialism requires.

In section 3, I critically examine Sosa's proposal, which over-emphasizes the similarities between belief and suspension, resulting in an over-intellectualization of these attitudes. I then develop my account as an extension of the knowledge-first virtue epistemology (KVE) that I develop elsewhere (Miracchi 2015*a*, Forthcoming), and argue that it is more plausible than Sosa's view. Whereas on Sosa's account epistemically assessable suspension is always a kind of epistemic performance with a complex aim (to believe if and only if doing so would be apt), on my view suspension is often a kind of non-performance that is derivatively rationally assessable as a part of our general epistemic comportment.

According to my account, beliefs constitutively aim at knowledge. They are *the* candidates for knowledge, and thus they (and only they) are the fundamental epistemic performances and so the primary bearers of epistemic evaluation. A belief is rational just in case in aiming at knowledge the believer is appropriately sensitive to what it takes to know on that question. The believer in such a case *manifests proper practical respect* for the constitutive aim of the performance. Such respect is compatible with failure, but only unlucky failure.

Suspension does not aim at knowledge (constitutively or otherwise), and so cannot manifest this kind of practical respect for what it takes to know. However, in suspending

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

one can manifest a kind of derivative proper practical respect. Precisely because it is a *withholding* of aiming at knowledge, when done competently it is a case of being appropriately sensitive to what it takes to know.

This account can explain why the same sort of assessment is applicable to both belief and suspension, as well as other epistemic activities like evidence gathering, etc. Epistemic rationality is the manifestation of proper practical respect for what it takes to know. Nevertheless, it also shows why belief has a privileged claim to these properties. Only belief can manifest this sort of respect toward its own constitutive aim. I close by showing how this account of suspension provides a more compelling response to Pyrrhonian skepticism than Sosa's account.

A note before proceeding: Throughout I will be largely unconcerned with the epistemologist's distinction between rationality and justification, which roughly correspond to internalist and externalist epistemic properties. The account I develop will show how suspension can have both kinds of properties. So, unless otherwise noted, please do not read "rational" and "justified" as tracking one of these notions instead of another.

1. The rational assessability – and possible primacy – of suspension

We regularly assess people not just for what they believe, but also for what they do not believe. Sometimes these assessments are negative: we think that those who do not believe in human-caused climate change are failing to believe what is required by the available scientific evidence. Sometimes these assessments are positive: we often praise people for not jumping to conclusions, and for keeping an open mind until it is clear what to believe.

Withholding belief comes in (at least) four kinds, and they are all epistemically assessable. First, there is the kind of suspension that comes from careful reflection. This is a settled attitude, a closing of inquiry on a certain question when one takes there to be sufficient reason to think that one is permanently (or indefinitely) unable to answer a question. Someone who carefully arrives at agnosticism about God, for example, might be praised as having appropriately reflected on what we are capable of knowing in the relevant domains.[1] Second, there is the kind of suspension that motivates deliberate epistemic activities other than belief – gathering more evidence, reasoning more carefully through a question, exploring analogies, asking questions, and listening more carefully. Such suspension is active: it is not a settled attitude but a deliberate withholding in the context of searching for the truth or for understanding.

Third and fourth, there are the cases where we implicitly or automatically withhold belief, either by allocating our attention elsewhere without answering the question of interest (a kind of closing of inquiry analogous to the first kind of suspension), or by engaging in the kinds of epistemic activities that would help us to answer it (a way of pursuing inquiry, analogous to the second).

We hold people epistemically accountable for all these kinds of suspension, and I take it as given here that our practices are sometimes reflective of the epistemic facts. Suspension is sometimes rational and sometimes irrational. Often in academic epistemology we treat suspension as if all forms are normatively equivalent, though they are not. Sometimes it is crucial not only to withhold judgment on a question, but to also

[1] I am an atheist. However, I think agnosticism and theism are both rationally permissible positions. As the reader will see below, on my view whether a belief is rational depends on the epistemic competences that the individual brings to the question at hand, and these can vary substantially. This does not mean that anything goes, but more than enough for there to be rational disagreement on such a difficult question.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

continue pursuing it.[2] Sometimes questions are difficult enough that a reflective attitude is required (for those generally capable of adopting it). Still, it's often crucial to have and cultivate the right first-order dispositions – if we're not disposed to hesitate in the face of inadequate evidence, our reflective capacities won't be much help. Especially given the state of today's politics, we must be quick to exhibit doubt, examine our sources, and fact-check.[3] Still, a satisfactory account of suspension should explain what is normatively in common to all these cases, even if there are important differences.

Accepting that suspension is epistemically assessable raises many questions. What is the relationship between the kind of assessability appropriate to suspension and that of belief? Does one have a kind of primacy over another, or are they of the same kind? Many skeptics take suspension to have a kind of default or privileged epistemic status – in the face of doubts about the epistemic appropriateness of belief, the rational response is suspension (of one or another variety). Pyrrhonian skeptics, for example, plausibly take suspension to be the least risky of all possible options. Sosa attributes to them the following claim, which I will call *Skeptical Default*:

> **Skeptical Default.** Insufficient reason to judge (positively or negatively) provides sufficient reason to suspend (Miracchi Forthcoming: 1).

This attribution treats the Pyrrhonian as endorsing suspension as a deliberate, settled attitude. While some skeptics may do this (perhaps Descartes or Hume in their more skeptical moods), I do not think this is the best way of understanding the Pyrrhonian outlook. Sextus Empiricus, for example, is quite clear that he is describing a cultural *practice* of continuous inquiry (investigation of opposites), instead of a policy or commitment to suspend belief. Although this practice has the result that the epistemic agent does not form any settled opinions, this is a by-product:

> Since we have been saying that tranquility follows suspension of judgment about everything, it will be apposite here to say how suspension of judgment comes about for us. It comes about – to put it rather generally – through the opposition of things. We oppose what appears to what appears, or what is thought of to what is thought of, or crosswise. (Sextus Empiricus 2000: Bk. I Sec. xiii)

By keeping his account of the skeptical approach descriptive and focusing on practices rather than commitments, Sextus Empiricus (at least prima facie) avoids commitment to any claims, including the claim that one *should* suspend belief:

> Thus, although the phrase 'In no way more' exhibits the distinctive character of assent or denial, we do not use it in this way: we use it indifferently and in a loose sense, either for a question or for 'I do not know which of these things I should assent to and which not to assent to'. Our intention is to make clear what is apparent to us, and as to what phrase we can use to make this clear we are indifferent. Note too that when we utter the phrase 'In no way more' we are

[2]Some formal epistemology misses out on this by focusing exclusively on the rationality of credal updates. Where one should not believe that p, it is often much less important (and perhaps indeterminate) what degree of confidence one should have in a proposition and much more clear that one should be engaged in doing what one needs to do to figure it out.

[3]An interesting issue is whether reflective suspension (of either the first or second kind) would be possible without first-order suspension (the third kind). I think not, but cannot pursue this question here.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

> not affirming that it is itself certainly true and firm: here we too are only saying how things appear to us. (Sextus Empiricus 2000: Bk. I Sec. xix)

Our third and fourth kinds of suspension then – the automatic, unreflective, first-order kinds – are those most appropriately attributed to the Pyrrhonian skeptic. The fourth is attributed in the practice of opposing opposites, the third in the "tranquility" that results from this practice.

Why should these kinds of suspension be epistemically least risky? Although deliberate suspension is not itself a belief, it seems that coherence requires that if we *deliberately* suspend on a question then we should not reject certain meta-beliefs as irrational: for example the belief that one is rightly suspending or that one should suspend. Indeed, *Skeptical Default* is hard to make sense of unless it claims that such a belief is rational. Thus, if one is really skeptical about whether there is sufficient reason to believe *anything*, not just facts about one's cognitive access to the external world, for example, but any claim whatsoever, then deliberate suspension remains a risky option.[4]

So, we can characterize the Pyrrhonian as presenting apparently safe epistemic courses of action, mostly through example:

> **Skeptical Practices.** Whenever one claim appears to be true to us, we practice opposing the reasoning that motivates that claim with reasoning motivating its negation. In this way (as it happens) we avoid commitment to any claims, even that suspension is the safest epistemic route, although it is manifest in our practices that that is how things appear to us.

At this point it is appropriate to clarify the aim of the anti-skeptical project I am engaging in here (see also Miracchi 2017*b*). One need not try to convince the skeptic to abandon her ways in order to provide an illuminating anti-skeptical response. By revealing how certain commitments of our own epistemological theorizing might at first have seemed plausible but lead to skeptical results, we can better refine and develop our account of what epistemic relations there are and when we have them. When approached in this light, skeptical concerns are interesting because they reveal to us a failure in our theorizing about our epistemic relations in the actual world.

Adequately responding to the Pyrrhonian skeptic will not merely involve getting our reflective ducks in order; we will have to examine how the rational epistemic agent conducts herself generally, whether she reflects on or deliberately chooses this conduct or not. An adequate epistemology of suspension will explain what the rationality of suspension is such that it can't be adopted in this kind of blanket, universal way, even on the first order.

Attention to Pyrrhonian skepticism is interesting not only for academic reasons. If our epistemic theories always sanction suspension as the epistemically safest option, we cannot explain the impermissibility of *opting out* of certain doxastic practices: the irrationality of those who suspend on human-made climate change or systemic racism, for example. Moreover, many climate change deniers (etc.) may be using no determinate epistemic policies at all, and definitely not ones they explicitly defend. Attending to Pyrrhonian skepticism, then, may help us say something useful even when we cannot isolate any general policies or commitments that we can show to be irrational. This project is useful even if climate change deniers are not themselves convinced by our

[4]This does not make it any less interesting to argue against those who think that deliberative suspension has the epistemic upper hand – just that one's target should not be the Pyrrhonian.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

response: there are onlookers in these debates, and the public discussion can be significantly changed by demonstrating the epistemic irrationality of such positions.

Attention to the case of Pyrrhonian skepticism is thus not merely academic: it can help us respond in contemporary public debates to those who present themselves as taking the rational, careful, upper hand when doing so is deeply irresponsible and even disastrous. It is important not to cede epistemic ground in these debates.

Sosa, then, does not include the whole problem within his purview when he writes:

> Yet there must be something importantly in common between how *judging and suspending* are respectively justified. What needs to be justified is after all the *choice* between suspending, on one hand, and affirming, on the other (positively or negatively), on the given question. (Sosa Forthcoming: 3) (his emphasis)

For Sosa, what makes suspending rationally assessable, at least at first pass, is the connection it bears to deliberately held belief. When one chooses to believe, one chooses not to suspend. When one chooses to suspend, one chooses not to believe. Thus, insofar as the choice to believe is rational or irrational, so correlatively is the choice to suspend. However, we have already seen that focusing on deliberate suspension is too narrow, both in general and for the anti-Pyrrhonian project. The Pyrrhonian project is particularly interesting precisely because it guides us in developing an epistemological theory that does not generally license suspension of any kind, even on the first-order.

To see how difficult this kind of anti-skeptical project is, let us examine how one popular contemporary view fails to do this. If we analyze epistemic normative properties as kinds of norm compliance or violation, as for example Timothy Williamson (Forthcoming) and Clayton Littlejohn (2012) do, we will have difficulty countering the Pyrrhonian skeptic. Williamson is sensitive to this when he writes:

> A Pyrrhonist sceptic may hope to comply vacuously with all three norms [(N) Believe only what you know, (DN) Be the sort of person who is disposed to believe only what one knows, (ODN) Do the thing that a person who is disposed to believe only what she knows would do] by having a general disposition never to believe anything. If one has no beliefs, then a fortiori one has no untrue beliefs, no beliefs that fail to constitute knowledge, no beliefs that are improbable on one's evidence, no inconsistent beliefs, and so on. The Pyrrhonist, if such a person is possible, complies with all three norms even in the sceptical scenario. … Non-sceptics may find little to admire in the Pyrrhonist's self-imposed ignorance, especially when that ignorance concerns the needs of others. There may be positive norms for knowledge, such as a norm enjoining knowledge-gathering in various circumstances, and so positive as well as negative norms for beliefs. (Williamson Forthcoming)

Here Williamson acknowledges that his account of the normative properties relevant to knowledge countenances the radical behavior of the Pyrrhonian skeptic as rational, because it is a kind of vacuous norm compliance. Indeed, being a Pyrrhonian is actually the surest way to comply with epistemic norms. He remarks that one may not like this result, and so impose further constraints to rule it out, but there is nothing in his account of rational belief that rules out the Pyrrhonian's practices.

There are two reasons why the Williamsonian anti-skeptical strategy of imposing further constraints on epistemic attitudes is inadequate. First, it is ad hoc. Rather than being motivated by an account of what rational or justified epistemic comportment is, such constraints are motivated by a desire not to let the Pyrrhonian have the upper hand. Our epistemological theory should provide a motivated reason for

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

why the Pyrrhonian is at epistemic fault. Second, it fails to illuminate the connection, noted above by me and Sosa, between rational suspension and rational belief. One wants an account that explains not only why suspension is permissible in cases where belief is not, but why suspension (of one kind or another) is often impermissible in cases where belief is permissible.

To summarize thus far, we can outline three desiderata for a theory of the rational assessability of suspension, besides general extensional adequacy:

Desiderata for an account of the rational assessability of suspension:

1. The theory assesses some cases of suspension as rational and some as irrational.
2. The theory rationally assesses first-order suspension as well as deliberate suspension (of either the settled or inquisitive variety).
3. The theory accomplishes (1) and (2) as part of a general account of the rationality of epistemically assessable attitudes, including belief.

These are the desiderata that I will take to be required for an account of suspension in what follows. By acknowledging the epistemic assessability of suspension and the need to provide a unified account of rational suspension and rational belief, one may be tempted to include suspension among the kinds of attitudes that are *fundamentally* epistemically assessable, i.e. on par with beliefs. Let us call this commitment *Rational Parity*:

Rational Parity. The features that make epistemic rational assessment applicable to suspension are the very same features that make such assessment applicable to beliefs.

This is Sosa's (Forthcoming) strategy, as we'll see below. It is attractive because it helps us conceive of suspension positively. Suspension will be a kind of positive epistemic attitude that is subject to epistemic requirements, just as belief is. This approach thus precludes thinking of suspension as a kind of *opting out* that can vacuously satisfy epistemic norms, as we saw was problematic in Williamson's case.

As attractive as this approach is, I think it should be rejected. I will argue for this in section 3, but I'll merely note here the thought that opened this paper, namely that the rational assessability of beliefs seems fundamental to epistemology, and is often treated as such, whereas suspension does not have such a central place. Thus, it may be preferable, insofar as we can give a general and unified account of the rationality of epistemic attitudes, to treat the normativity of suspension as derivative from that of belief:

Rational Derivativeness. The features that make epistemic rational assessment applicable to suspension are derivative from the features that make such assessment applicable to beliefs.

I shall explain below how we can provide an account of suspension that endorses *Rational Derivativeness* and still satisfies the desiderata outlined above. First, however, as a way of better understanding what is required of such a theory, let us investigate the shortcomings of evidentialism on the issue.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

2. Challenges to evidentialism

For the purposes of this paper we can characterize evidentialism as follows.

> **Evidentialism:** *S* is doxastically justified in believing *p* if and only if (Ei) *S* bases *p* on *E*.
> (Eii) *E* is S's total body of relevant evidence.
> (Eiii) An adequate "evidencing" relation holds between the propositions in *E* and *p*.

This characterization of evidentialism does not depend on one's conception of evidence, beyond that it is propositional. One's evidence could be internal mental states, knowledge, justified belief, credences, subjective or objective reasons. What is important for our purposes is that evidentialism claims that we can analyze doxastic justification in terms of the obtaining of an adequate propositional connection between *E* and *p*, and an adequate sensitivity to this connection. Sosa uses the term "evidences" here (Forthcoming: 2), which is helpful because it points out the possibility that one's evidence might not support the taking of an attitude even if certain evidential relations hold between it and the attitude's content. Let us throughout assume whatever account of the requisite propositional connection between *E* and *p* the evidentialist likes, and say that in such a case *E* "evidences" *p*, reserving the term "support" for the appropriate doxastic connection.

Sosa's first claim is that evidentialism so understood must be augmented by further conditions on the basing relation. It's not enough that *S* bases *p* on *E*. *S* must do so because *S* is sensitive to the fact that (Eii) and (Eiii) obtain. So we should emend (Ei) on behalf of the evidentialist as follows:

> (Ei*) *S* bases *p* on *E* in a way that involves taking (Eii) and (Eiii) to be satisfied as the rational motivation for basing *p* on *E*.

This kind of modification is within the spirit of evidentialism because it still distinguishes whether one's evidence adequately supports *p* from psychological or other individualistic features about how the agent epistemically comports herself. For the evidentialist, what makes *E* sufficient evidence for believing *p* depends only on facts about the propositions in *E* and *p*, not further psychological facts about the agent. Proper epistemological comportment is exhausted by proper grasp of such evidential relations and use of these relations in forming beliefs. It is this idea that is Sosa's (Forthcoming) main target in arguing against the evidentialist, and rightly so. As we shall see, epistemically assessable comportment extends much further than the appreciation of evidential relations, and this makes a crucial impact on the rationality of both suspension and belief.

2.1. Sosa's first challenge

Sosa's first challenge is as follows:

> It is hard to see how *suspending judgment* on a given question < p?> could ever be *based* on one's total evidence. *Judging* can plausibly be based on evidence if the evidence speaks sufficiently in favor of the content of your affirmation. But it is unclear how suspending could be *based on evidence* in virtue of some relation between the evidence and the content of your suspending. Sosa (Forthcoming: 3)

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

I take it that the worry is as follows. Given evidentialism's account of justification in terms of basing, one cannot be justified in suspending. For *ex hypothesi* the kind of basing necessary for justification only obtains when *E* evidences *p*. So there is no such thing as properly basing suspension on *E*. So there is no such thing as epistemically rational suspension.

This challenge assumes *Rational Parity*, as characterized above. Sosa is assuming that suspension is rational *in the very same way* that belief is rational. If so, then the basing relation would have to be the same for both suspension and belief. However, it is open to the evidentialist to deny this, accepting *Rational Derivativeness* instead. For example, one could hold that rationally suspending on *p* is what you do when (Eii) is satisfied and (Eiii) is not, and you take this to be the case (implicitly or explicitly). When one suspends on *p* based on *E* in this way, *E* supports one's suspension differently. This is a natural extension of the account. It plausibly delivers a sense in which suspending is genuinely epistemically rational: the subject withholds because she is appropriately sensitive to the propositional relationship between *E* and *p* and whether it is an evidencing relationship. What attitude is supported by *E* depends on proper grasp of this relationship. In this way, the evidentialist's account is general and unified.

It is further open to such an evidentialist to claim that in cases where *E* does sufficiently evidence *p*, suspension is *irrational*. Adopting this commitment would provide an account of the (ir)rationality of suspension that explains how it is derivative from that of belief and that rules out Pyrrhonian skepticism as the rational default. On this kind of approach, the evidentialist can claim that taking on any rational epistemic attitude involves believing with proper sensitivity to the evidential support relations. Sometimes these rationalize belief; in all other cases they rationalize suspension. This is an intuitive way of characterizing *proper basing* in a broader way, one that plausibly captures what is correct about Sosa's claim that what is rational or justified is the "choice" between believing and suspending.[5]

So, Sosa's first challenge can be satisfactorily resisted by the evidentialist. Let us examine his second challenge, which I think is deeply correct and does pose an insurmountable problem for the evidentialist in a way that favors a virtue epistemological approach.

2.2. Sosa's second challenge

This second challenge derives from the observation that rational suspension requires that the epistemic agent be adequately sensitive to whether her total evidence is sufficient for believing *p*. The challenge is that, contra the evidentialist's claims, whether this is the case is typically not a question that *E* itself properly answers, but instead has to do with how the subject acquired her evidence, what alternatives are available, and what she is disposed to do with other bodies of evidence. If whether *E* is sufficient support for belief in *p* is not settled by the contents of *E*, then evidentialism will not be able to provide an adequate account of rational suspension:

> One must avoid the negligence or recklessness of judging on a body of evidence prematurely. Such *negligence* or *recklessness* would preclude creditably competent performance. … [For example] To insist on mental arithmetic, while willfully

[5]One potential drawback of this approach is that it does not countenance certain kinds of risk aversion as permissible when they intuitively are: in some cases where it is permissible to believe it also seems permissible to suspend.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

> ignoring easily available, more reliable methods, is to incur epistemic negligence or recklessness. (Sosa Forthcoming: 6)

We can formulate the crucial issue as follows. *E* might evidence *p* – the requisite propositional connections might hold – but only because *E* is a very partial and misleading subset of the relevant evidence that's "out there". In such cases, one is not justified in believing *p* on *E*. For example, consider what has become known as "the Republican Memo" (or "the Nunes Memo"), which alleges that the FBI inappropriately surveilled the 2016 Trump campaign, and included the public release of heavily redacted classified information to support this claim.[6] Plausibly the information provided in the memo evidences the claim that the Trump campaign was inappropriately surveilled. However, belief in this claim is not warranted by the memo precisely because it does not include enough relevant information and is likely to be only a very skewed subset. At most, suspension is warranted, but so plausibly is outright disbelief.

Of course, according to the evidentialist whether a subject rationally believes that the FBI inappropriately surveilled the Trump campaign depends on their *total* evidence, not just the memo. Such a subject plausibly would not have as part of their evidence that the Republican Memo presents a very skewed presentation of the relevant information. Nevertheless, it is highly plausible that because there is such evidence "in the offing", so to speak, epistemic agents should be held responsible for not having it, and so for not suspending in this case.

Similarly, when *E* fails to evidence *p*, not just any form of suspension is permissible – for example, it may be impermissible to adopt an attitude of settled suspension. Perhaps further evidence-gathering is required, and perhaps in the meantime belief in *p*. For example, it is plausible that an academic who hears rumors about a colleague's impropriety with female graduate students is not epistemically justified in settling on suspension (type 1 or 3 characterized in section 1). There are too many cases where hearing such rumors is a precursor to conclusive positive evidence. (This concern is distinct from whatever *ethical* reasons one might have not to settle on suspension. I will return to this below.) Instead, the subject is at least epistemically required to gather more evidence, and may even in the meantime be required to adopt belief in – or at least acceptance of – their colleague's guilt. This is true, I submit, even in cases where the epistemic agent does not have as part of their evidence that hearing such rumors *is* a precursor to conclusive positive evidence. Because there is such evidence available, epistemic agents can be held responsible for *not* having it.

What is the upshot here? An account of what it is to properly take one's evidence to support a conclusion often depends on factors that go beyond the propositions that are in one's evidence. In other words, the mere "taking" condition (Ei*) I offered on behalf of evidentialism above isn't enough. Suspending and believing properly on the basis of evidence both require more than appreciating the propositional connections between your evidence and the candidate proposition for belief. They require you to have the right evidence, to have the right dispositions to gather more evidence when you need it, and to not rest or settle with evidence that does not properly settle your questions. This takes us outside of the realm of evidentialism proper, into dealing with responsibility for one's epistemic comportment more generally.

2.3. Objections and responses

This Sosa-inspired argument just given is fairly quick, but I think it holds up. Here I'll examine some objections and replies to show why.

[6]Bump (2018), NPR (2018).

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

Objection #1: Subjects really are rational in these cases, because all it takes to be rational is to believe or suspend in accordance with your evidence.

Whether this line is motivated by intuition or theory, it should be rejected. First, although intuitions are often an important guide in epistemological theorizing, the intuitions here are too varied, involving well-worn differences between those with internalist- and externalist-leaning intuitions. Instead, the claim that we should restrict our primary epistemic assessments of one another (which are marked by the terms "rational" and "justified") to how one deals with one's evidence should be theoretically motivated.

A crucial, and I think non-negotiable role for our primary epistemic assessments is that we use them to hold each other accountable for what we believe or do not believe, and why. It is hard to see what point epistemic assessments would have if they did not have this role, or why epistemologists should be so interested in studying them. There may be a positive normative property possessed by people who believe in accordance with their evidence no matter how that evidence was arrived at, but such a property does not deserve a primary place in our epistemological theorizing. This is because we in fact *do* have substantial control over what evidence we have, what reasoning capabilities we have, whether we jump to conclusions, and so on. This control is often diachronic, and so does not entail any kind of problematic voluntarism about belief. Because we do, over time, have control over the quality of our bodies of evidence, such quality can properly figure into how we hold people accountable for their epistemic attitudes.[7]

Moreover, part of the point of epistemically assessing others is not just descriptive, but prescriptive. By assessing others as believing properly or not (rationally or not, justifiedly or not), we are effectively either sanctioning their methods of epistemic comportment or we are urging them to change. Because we do, over time, have control over the quality of our bodies of evidence, it is often *important* that we not sanction poor practices of gathering evidence, even if the person was not in control over the development of such poor practices. In making epistemic assessments, we should encourage each other to believe and suspend in better ways going forward.

Objection #2: There are general evidential facts that go beyond what evidence the agent actually possesses that we can use to make the necessary epistemic distinctions. For example, the rationality of belief or suspension may be a function of how the agent's body of evidence compares to the *total possible* evidence relevant to the issue at hand.

This approach tries to save the evidentialist by including facts beyond the subject's actual evidence, but it's too strong. Just because some evidence is "out there" doesn't mean one is responsible for having it. It might be too onerous to acquire, or might need some special creativity to be discovered. And, just because one person does discover such evidence doesn't automatically make it the case that others become irrational or unjustified in their positions. Instead, subtle factors are at play in determining whether a person is epistemically responsible for having more, or different, evidence than she actually does. It's not merely whether the evidence the subject has is partial or misleading, but whether she has done what is required of her in the evidence-gathering process. In order for this strategy to work, evidentialists would need to give us an account of when this is the case.

[7]For similar work along this line, see Peels (2017).

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

However, I think that it is unlikely that they will be able to do so in a way that retains the spirit of evidentialism. This is because it is not merely the propositions of our potential or actual evidence, but also facts about our mental capacities, life experiences, and other psychological features that make an epistemic difference. For example, there are cases of genuine rational disagreement where parties to the debate have different bodies of evidence and neither has failed to do their due diligence in gathering evidence. Academic disagreements are often like this: it is impossible to read everything relevant to an issue, and so researchers with different interests and proclivities will have different bodies of evidence. This may require them to adopt different epistemic attitudes, including suspension. Plausibly, in such cases often one if not both bodies of evidence are highly partial with respect to the total possible evidence on a question, and one or both are likely to be misleading. Still, it might be onerous or perhaps impossible to find out whether and how this is the case. We should not hold all academic commitments to be irrational or unjustified in such cases.[8]

Instead of trying to fit our wider range of epistemic practices into some purportedly sufficient abstract propositional relation between bodies of evidence we actually or potentially possess and propositions up for debate, we should work to understand what it is to do one's *due diligence* in gathering and maintaining one's evidence base. This again takes us outside of the realm of evidentialism proper.

> *Objection #3*: Although the "evidencing" relation won't be sufficient to distinguish the cases of interest, adding extra conditions to the *basing relation* can account for the differences.

The idea of this strategy is to build sensitivity into whether E provides sufficient support for p into the (proper) basing relation. However, we cannot adequately do this and still retain the spirit of evidentialism. This is because whether E is really enough for settling p? isn't something that can be answered generally, for any E and p. It depends on the particular contents of E, what p is, who the epistemic agent is, and contingent facts about the world. To see this, consider the following case:

> **Ignorance from cultivated incompetence.** Mike is a man who throughout his life has been excused from emotional and other kinds of service labor due to perceived brilliance. Due to this freed up time and energy, he has focused on his academic pursuits to much success. He is implicitly aware that not spending time on such labor promotes his academic success, and so avoids situations and responsibilities where he would learn how to do these jobs. As such, he is further excused from them because he is not currently capable of performing them. Mike therefore has little evidence about the emotional and service labors that his other colleagues are performing. For example, his evidence from his own experience is that students rarely ask for meetings outside of office hours, and that when they do it is easy and costless to refuse to meet. He thus refuses such meetings in order to focus on his research. He believes based on this experience and other similar evidence that his female colleague who regularly meets with students outside of her office hours does not value her research as much as he does.

Unfortunately, this sort of phenomenon will be familiar to many. By using the term "cultivated incompetence", I mean to draw attention to the long-term agency we have

[8]See Will Fleisher (2017) for a different approach, on which academics should not believe their professional views, but rather take a different kind of attitude (*endorsement*) towards them.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

over what we are competent or incompetent at, and so the responsibility we bear for certain failures even if in the moment we are incapable of doing better. Even though Mike's development of his own incompetence in service, organizational, and emotional labor was not deliberately planned, was bound up in the context of other concerns, and was socially sanctioned and supported, he still has enough awareness to realize that continued avoidance of and incompetence in this work is to his advantage, and that he has control over his level of competence in such work.

Cultivated incompetence is epistemically interesting (at least) in that it produces systematic ignorance on related topics: both the impoverishment of evidence and the incapacity to reason appropriately from that evidence (perhaps through lack of practice).[9] Here I focus on the former impoverishment, abstracting from the latter.

The way I have set up the case, Mike need have no evidence about this epistemic impoverishment. It is possible (and not unrealistic) to suppose that he does not spend enough time with his colleagues who regularly do this sort of labor to acquire the relevant evidence. I take it that there's a plausible way of filling out the case where Mike's evidence really does support the claim that his female colleague is making different autonomous choices about her time and so doesn't value her research as much as he does. However, given how widely available knowledge about systemic sexism and racism (etc.) in the academy is to professors, including knowledge about the extra emotional labor and time women professors are regularly expected to take with their students, *Mike should know better* than to believe on such evidence.

The way Mike forms beliefs is out of touch with what it takes to get onto the facts in this domain, and it is not too onerous to require him to expand his evidence base in a way that would substantially change the evidencing relations for similar propositions. He should suspend: he should either keep an open mind on the matter, or go gather more relevant evidence, or drop the question altogether. In claiming that Mike's belief is unjustified, we are properly holding Mike responsible for having this impoverished evidence base and criticizing his epistemic comportment on a particular occasion. Our epistemic evaluations – of subjects and particular beliefs – thus are sensitive to broader features of our agency than how we respond to the evidence we in fact have. An adequate theory of rational belief and suspension must take these broader features of our epistemic agency into account.

However, no general modification of the basing relation will be able to account for Mike's case in a way that appropriately distinguishes it from other cases. This is because it is particular facts about misogyny that make Mike's evidence base too impoverished. Compare, for example, a similar judgment made by a devout yoga practitioner, who judges that their neighbor prefers tennis to yoga because they frequently see them going to play tennis but not yoga practice. It seems rational for her to conclude that her neighbor prefers tennis to yoga.[10] Whatever the "evidencing" relations are, it is plausible that sufficiently similar relations occur in the research/teaching case as in the yoga/tennis case. Yet the former is not rational and the latter is.

Furthermore, the facts that make Mike's beliefs unjustified are contingent. Let us hope that there's another possible world without sexism and so Mike's evidence base really does provide sufficient evidence for his beliefs about his colleague. So in another situation those same evidencing relations *would* provide sufficient support for belief.

[9]See Briana Toole (2018) for an excellent defense of the claim that our social identity makes a difference to both our evidence and our competence in reasoning with evidence.

[10]Note that even here, however, it matters that tennis and yoga are both physical activities that typically require fairly high socio-economic status for consistent access. Where these socioeconomic facts differ, such suppositions will also be problematic.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

After all, it's not always problematic to generalize from one's experience to others' preferences (cf. the yoga/tennis case), and it is too burdensome to require that we always have knowledge or evidence of when it is appropriate. Small children, for example, have lots of knowledge about others' preferences, even if they do not or cannot have knowledge of when such inferences are permissible.

Moreover, note that the objections presented here do not depend on an internalist account of evidence, or indeed any account of what is required for propositions to be part of one's evidence. For example, we might suppose that your evidence is what you know (E = K), as advocated by Williamson (2000), and still get the same result. What makes Mike's evidence base impoverished has nothing to do with the quality of the individual pieces of evidence, but rather how the modes in which he acquires evidence skew the total body of evidence in ways that systematically lead him astray. Evidentialism has nothing to say about this problem.

Even very sophisticated versions of evidentialism that incorporate knowledge-first and virtue-theoretic components have difficulty with this case. For example, on Errol Lord's (2018) view, one only rationally believes that *p* if one believes that *p* for an objective (factive) reason *r* to do so, and manifests know-how to use *r* as a reason to believe *p*. Lord might hold either that Mike doesn't have sufficient objective reason to believe that his colleague values her research less than he does, or that he fails to manifest knowledge of how to properly believe this claim based on the objective reasons he has.[11]

However, things are not so simple. Lord, like other evidentialists, rejects the idea that reasons you do not possess can affect what is rational for you to do. He spends a good amount of space in the book showing how one can vindicate the intuition that it can be rational to do things even if there are objective reasons not to; for example, it can be rational to eat fish that contains salmonella if one is unaware of that fact (p. 23). Presumably in such a case it is also rational to believe that it is okay to eat the fish. If that's true, then Lord and other evidentialists owe us an account of what the difference is between the salmonella case and Mike's case. Mike has a lot of evidence that, were contingent facts about sexism different, would be sufficient to rationalize his belief. Why are these not sufficient objective reasons? Why does he not manifest his know-how? Lord must explain why Mike's apparent reasons are actually not objective reasons in the actual world or why he fails to properly base his belief on his reasons.

I think the simpler and more plausible explanation of the difference is that Mike is responsible for having more evidence than he does, while the subject in the salmonella case isn't (she's in a restaurant with the appropriate certifications, etc.). Where we have significant diachronic control over possessing highly relevant evidence, it would not be too onerous to do so, and we have failed to exercise such control, then we can be properly held responsible for failing to possess it.[12]

I don't deny that there is *something* positive about Mike's epistemic situation, or generally believing in accordance with one's evidence. What is at issue here is whether this kind of epistemic status deserves a central place in our theorizing or deserves the name "justification" or "rationality", which are tied up with the concepts of responsibility, praise, and blame. Once we start to pay more attention to how we have diachronic control over what evidence we have and how we acquire and maintain it, it becomes quite plausible that these factors are relevant to the epistemic evaluations that matter most.

[11]Thanks to Kurt Sylvan for encouraging me to address this question.

[12]I will return to exactly what the lesson from Mike's case should be below. Here I am focused on how evidentialism cannot account for such cases as a matter of the structural features of the view.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

Objection #4: The concerns discussed above about evidentialism are merely ethical concerns, not properly epistemic concerns. Perhaps holding one another accountable for our epistemic attitudes in a broader sense depends on all these features, but when we consider assessment that is properly epistemic, the evidential support relations do not divorce (so strongly at least) from the evidential support relations.[13]

Most of the examples discussed above do indeed involve important ethical concerns, but the normative assessments I am discussing here are genuinely epistemic. First, it may be that our epistemic assessments are not wholly divorceable from ethical concerns. If some version of socially (as opposed to individually) determined pragmatic encroachment is true, then there may be higher standards for justification in ethically charged cases.[14] However, although the approach advocated here is compatible with the possibility that pragmatic factors might change the standards required for justification, none of the examples or arguments presented here rely on this.

Instead, I use ethically charged examples because these are examples where we can easily see the importance of a key role of our epistemic concepts: that of holding each other accountable for what we believe and why. It is perfectly reasonable to read this as holding each other *epistemically* accountable. Although we are (epistemically) permitted to hold each other epistemically accountable on any question, on issues where a person's beliefs do not make a significant pragmatic or ethical difference we will often be less willing to "rock the boat" by criticizing one another's epistemic comportment.

Indeed, I think that the attractiveness of evidentialism largely rests on restricting our purview to cases where what we believe and why do not have substantial effects (brain-in-a-vat scenarios being the extreme case. One might think that in doing so one is ruling out irrelevant factors so that one can more clearly focus on the genuinely epistemic facts. However, there is good reason to think this is mistaken. If harm is indeed a dimension that affects our epistemic assessments – whether correctly or not – there is no reason to think that the harm*less* case is somehow the neutral case, where we get at the true epistemic assessments. An exclusive focus on such cases may be leading us astray just as much as the opposite focus, if not more so. We must be careful throughout to tease apart genuinely epistemic normative assessments from other kinds of assessment and to supplement our intuitions with arguments.

In all of the examples given above (involving the Republican Memo, graduate misconduct, academic disagreement, and cultivated incompetence) we plausibly hold the subjects epistemically responsible because their evidence bases are too impoverished to reliably and systematically get them onto the facts in the domains in question, even though it is within their control to do something about it. In all these cases the subjects' evidential states are under their (diachronic) control, from the past to the moment of assessment and moving forward. Our judgments indicate that our epistemic assessments of one another are sensitive to our broader epistemic agency, extending to our evidence gathering and handling practices generally (and beyond). These are properly epistemic assessments because they are concerned with what it takes to get onto the

[13]Thanks to Timothy Williamson for encouraging me to discuss this concern.

[14]Although this view is yet to be defended in the literature, it aligns with some work in social epistemology, e.g. by Sandy Goldberg (2018). On such a view, it wouldn't be just one's personal stakes that affected the standards for justification, but rather societal stakes. Someone who is attracted to the individual pragmatic encroachment thesis should plausibly be attracted to this social thesis, given the insights from the social epistemology literature about the legitimate epistemic expectations we have of one another.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

facts. The evidentialist therefore cannot dismiss these arguments as showing that believing in accordance with one's evidence is merely ethically insufficient, as opposed to epistemically insufficient. Truth-related factors and concerns outside of one's evidence are importantly relevant to our proper epistemic assessments of one another.

This relates to a point made in response to the first objection I considered above. Our epistemic assessments of one another do, and should, make a difference to how we epistemically comport ourselves. By not taking into account how our broader epistemic practices affect one another, we are effectively condoning lazy or otherwise insufficient ones. Sosa's key insight, which I have expanded on here (in ways he might or might not endorse) is that appropriate epistemic assessment legitimately takes into account the full range of epistemic powers possessed by the assessed agent, including practices of evidence gathering, the development of certain patterns of reasoning and thinking, the cultivation of new cognitive capacities, etc. When we fail to include all of these aspects of epistemic agency as relevant to our epistemic assessment, we are effectively letting people *off the hook* for being reckless or negligent in some of these aspects of their agency. There is no neutral position here.

I conclude that, in order to account for cases where suspension is required or inappropriate, we must move beyond evidentialism and embrace an account that focuses on the wider properties of the believer as an epistemic agent. We need a theory that takes into account not just how the believer treats the resources they have at the moment, but whether they have taken appropriate epistemic care to put themselves in a position to act aptly in the situations where they tend to find themselves.

3. A knowledge-first virtue epistemological account of suspension

In this section, I develop my knowledge-first virtue epistemological account of suspension through critical comparison with Sosa's proposal (2019, Forthcoming). This view extends the account of justified and rational belief I develop elsewhere (Miracchi 2015*a*, Forthcoming) to the case of suspension. I will argue that my view is simpler and more plausible than Sosa's view, and that it can provide a more compelling response to Pyrrhonian skepticism.

First, however, it's important to explain what the accounts have in common as versions of reliabilist virtue epistemology, and why we should expect virtue epistemology as a general approach to provide us with a better account of suspension than evidentialism. According to virtue epistemology, the epistemic domain is a performance domain, structured around epistemic aims and values and what it takes for an agent to achieve or manifest them. Reliabilist virtue epistemology puts a reliability condition on competence possession, so that one possesses a competence to attain an epistemic aim only if one reliably attains that aim when one exercises one's competence. A belief is justified or rational just in case it is an exercise of such an epistemic competence.

How such an account deals with suspension will depend on how the account is filled out. For example, what one takes to be the distinctively epistemic aims that structure the realm of epistemic normativity will make a significant difference. Traditionally for reliabilist virtue epistemologists, the fundamental epistemic value is true belief and the fundamental aim its attainment. As a knowledge firster, I claim that knowledge has these properties. On both approaches, exercises of epistemic competence cannot be cases of suspension because in suspending we do not attempt to believe truly, or to know. Indeed, suspension *is* a kind of non-believing, a not settling on an answer to the question of whether *p*. So, reliabilist virtue epistemology so far does not explain the epistemic assessability of suspension.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

Still, even at this level of generality, it should be clear why the approach might seem promising. Because epistemic normativity is understood in terms of the agency that a subject has in attaining what is epistemically valuable, it is the sort of approach that has the inherent ability to include features of our agency that go beyond, for example, how we deal with the evidence we happen to have in the moment.

How well virtue epistemology does in accounting for the broader factors relevant to our epistemic comportment will depend on the particulars of the view – for example how competences are typed. To see this, consider the (highly unrealistic) case in which Mike doesn't treat the differences between the way he forms beliefs about someone's research versus teaching preferences and many other cases of relative preferences where the way he forms beliefs is indeed reliable (as in the tennis vs. yoga case). The view I previously defended in Miracchi (2017*a*) entails that Mike is justified in his belief about his colleague.[15] However, if the arguments above are correct, the view must be revised. There are more external features that can rightly make a difference to our epistemic assessments – for example, it is not just relevant whether Mike reliably forms beliefs, but also that he could easily have done better with respect to a target domain. This latter fact can license a judgment of epistemic irrationality, even if the belief is reliably formed.

While exploring adequate criteria for typing epistemic competences is a project for another essay, I bring it up briefly here to illustrate that virtue epistemology does not *by itself* provide us with an account that is adequately sensitive to all the plausibly relevant features of our broader epistemic comportment. I now turn to discussion of Sosa's proposal and my own view.

As noted above, if the aim that explains epistemic rationality is true belief, suspension cannot be rational in the very same way that beliefs are, because it can't be an exercise of an epistemic competence to attain that aim. An account of suspension that satisfies *Rational Parity*, then, must offer a different epistemic aim to explain rationality. For this reason, Sosa (2019) claims that both belief and suspension are performances with the following aim:

> affirming alethically iff that would be apt (and otherwise suspend). (Sosa 2019)

By "alethically", he means "in the endeavor to affirm truly".[16] Note that the biconditional here allows for two ways for performances to satisfy this aim. One is to affirm when apt (belief). Another is not to affirm when affirming would not be apt (suspension). So both belief and suspension can count as performances with this same biconditionally specified aim. They can thus both be *competent* with respect to the same aim, and so rational in the same sense:

> Sometimes the exercise of an epistemic competence would yield judgment, but sometimes it would yield suspension. And both of these can be creditable epistemic attainments. (Sosa Forthcoming: 11)

While this account does bring suspension within the purview of epistemic agency, it generates some new problems. First, I think this account over-intellectualizes the aim of

[15]There I argue that the reliabilist (of any stripe) should type belief formation and maintenance processes by looking at how the believer treats these processes. These processes are of the same epistemic type when the believer treats them as such in determining her degree of confidence in the belief content.

[16]Because the aim is still truth-related, it has claim to being an epistemic aim (as opposed to, say, a pragmatic or ethical aim).

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

belief. The aim of belief is knowledge (or true belief, if you must). In believing, you aren't aiming to avoid error; you're aiming to know.

It's perhaps unfair for one virtue epistemologist to lob Aristotle quotes at another, but here I can't resist:

> All [humans] by nature desire to know. An indication of this is the delight we take in our senses; for even apart from their usefulness they are loved for themselves; and above all others the sense of sight. For not only with a view to action, but even when we are not going to do anything, we prefer seeing (one might say) to everything else. (Aristotle 1984: Metaphysics I.1)

More seriously, however, the kind of logically complex aim that Sosa appeals to is not plausibly psychologically real, at least not in general. Aiming at knowledge is one of the most basic performances that cognitive beings are capable of: the dog aims to know whether its owner will leave, the child aims to know what she can see around the corner. It requires too much of a sophisticated stance on one's relationship to one's epistemic environment to hold that in believing as they do these agents are really engaging in a complex performance of: *believing if believing would be apt and not-believing if believing would not be apt.*

Sosa thinks that the aim of knowing requires adopting the logically complex biconditional aim (Forthcoming: 11), but this is highly implausible. It is integral to any performance with an aim that in so performing one aims to achieve it. This aim is, of course, incompatible with failure, and performing with an aim entails that one has certain dispositions: e.g. to be satisfied when one has achieved one's aim and disappointed when one fails. However, as is widely acknowledged, we do not believe or aim for all the logical consequences of what we believe or aim for. It is thus not necessarily the case that in aiming to know we are also always aiming to avoid failure.

Sosa is primarily concerned with the intentional, reflective case, which is perhaps why the more sophisticated biconditional aim seems adequate to him. Recall the quote discussed above (my emphasis):

> Yet there must be something importantly in common between how *judging and suspending* are respectively justified. What needs to be justified is after all the *choice* between suspending, on one hand, and judging (positively or negatively), on the other, concerning the given question. (Sosa Forthcoming: 3) (his emphasis)

This takes the agency that is rationally assessable up a level from what we have been so far discussing. Performances can be competent or incompetent without being chosen. The competence of *choices* is often a further matter, as Sosa takes advantage of in developing his three-tiered epistemology of animal knowledge, reflective knowledge, and knowing full well, where the latter is a matter of one's apt judgment that one would aptly believe that *p* guiding one's first-order judgment.[17] In line with Sosa's other work, we should understand the "choice" here as an apt judgment about whether to believe an answer to *p*? which aptly guides one's believing or suspending. Sosa's account of suspension primarily, if not exclusively, applies to the suspension analogue of knowing full well.

Regardless of whether this is plausible for more sophisticated forms of withholding, we need an account that can deal with much simpler cases. This is in part because we can and do epistemically assess one another for withholding, even where this is not plausibly a choice. Recall the person who withholds belief about human-caused climate change in the face of overwhelming evidence in favor. We can and should hold such

[17] Sosa (2011).

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

suspension to be irrational even if it is not reflectively chosen, even if it is an automatic response to the evidence against the agent's culture and background commitments.

Secondly, by construing the aim of belief and suspension in this biconditional way, Sosa commits to a deep symmetry between belief and suspension where there plausibly isn't one. *Error is a contrast concept*. The very idea of avoiding error presupposes that another aim is already in place, one with a more fundamental claim to be an aim of the agent. In contrast with the aim to know, which is plausibly simple and basic to our psychology, the aim to avoid error is derivative, plausibly one that only beings who have the ability to reflect on their own agency can adopt.[18]

Along this vein, we return to the point that inappropriate suspension is not as much a failure as false belief. We *do* tolerate risk aversion more than we do recklessness. Now, we must be careful not to have too high a tolerance to risk aversion – otherwise we will not be able to answer the Pyrrhonian – but that does not preclude us from pointing out an asymmetry between belief and suspension in this respect. It is part of what makes the Pyrrhonian skeptical position plausible that suspension is often less risky than belief. However, on Sosa's account this is mysterious. Since the epistemic aim that explains the rationality of both is a symmetrical biconditional, one would expect – and indeed the account demands – exactly the same constraints on rational suspension as for rational belief.

Moreover, insofar as providing an anti-skeptical response to the Pyrrhonian is one of our goals, we should be careful not to construe withholding as some kind of commitment or reflective choice, which would carry with it a rational injunction to endorse one's withholding. As noted before, the Pyrrhonist is plausibly not *intentionally* withholding belief in *p*. She is suspending on the first-order, as a function of her broader epistemic practices (e.g. to oppose opposites). We thus need a way of understanding what she is doing that doesn't amount to a choice, or any kind of reflective attitude.

In summary: Sosa does provide a unified account of the rational assessability of more sophisticated forms of suspension. It assesses some cases of suspension as rational, and some as irrational, and so satisfies Desideratum 1. It also does so as part of a general account of the rationality of epistemically assessable attitudes, and so satisfies Desideratum 3. Moreover, although the paper does not deal with competence individuation, nothing about the view rules out a broader account, so that competences to *affirm alethically if and only if doing so would be apt* depend on a broader, diachronic suite of features related to one's epistemic comportment.

However, the account cannot explain important asymmetries between belief and suspension, and it does not adequately account for suspension's more regular and simpler forms. It thus fails to satisfy Desideratum 2, and so does not adequately counter the Pyrrhonian skeptic. These drawbacks, as we have seen, largely stem from Sosa's commitment to *Rational Parity*. If we give up this assumption, perhaps we can provide a more comprehensive account that can still explain the sense in which suspension is genuinely epistemically rational as part of a general theory of the epistemic assessability of attitudes.

3.1. A knowledge-first virtue epistemological account of suspension

On the view I develop in Miracchi (Forthcoming), epistemology narrowly construed is the performance domain structured by knowledge as the primary value and

[18]Even though on Sosa's view having the aim of aptly believing entails the biconditional aim and is in that sense "subordinate" to the former, the explanation of the epistemic assessability of suspension is still the same as that of belief. The rationality of suspension is due to its being an exercise of epistemic competence, the same epistemic competence as beliefs that are candidates for knowing full well.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

achievement. This means that only candidates for knowledge are the primary bearers of epistemic value. Rational or justified belief just is belief that is a *good candidate* for knowledge. I won't defend this view here, but I'll explain it enough to show how it can be extended to provide a novel and plausible account of suspension, one which adequately addresses the concerns raised above for both Sosa and the evidentialist.

There are two co-extensive ways of being a good candidate for knowledge, which are both necessarily possessed by exercises of competences to know. First, there is a probabilistic sense. Because of the reliability condition on competence possession, exercises of competences to know are as such likely to be cases of knowledge. More pertinent to our current concerns, there is a perspectival sense in which exercises of competence are good candidates for knowledge. Beliefs that are good candidates for knowledge manifest what I call *proper practical respect* for their constitutive aim.[19]

By this I mean that, in believing as one does, the agent is behaving in a way that respects what it takes to achieve the very aim of that performance, i.e. knowledge. The idea is probably best elucidated with paradigm examples. Think of the case of knowledgeably inferring *q* from *p* and *If p then q*. In forming a belief in *q* in this way, you're properly respecting what it takes to know. The way in which you're compelled to *q* on the first order manifests respect for what it takes to get onto the facts.

This kind of practical respect very plausibly *must* be accounted for on the first order, on pain of running into Lewis Carroll-style worries. Achilles cannot compel the tortoise to knowledgeably infer *q* by adding more premises (more evidence). The tortoise must at some point be compelled to infer *q* because he aims to know, and his epistemic position is sufficient to do so. Moreover, although this case clearly has a phenomenological component (the tortoise either "gets" it or doesn't), it also has a factive component: getting it isn't merely seeming to get it. The person who believes *q* because they "see" that it follows from *p* and *If p then q* is perspectivally sensitive to what it takes to know on this question. This just is a case of an epistemic agent competently aiming at knowledge in believing as she does.[20]

Perhaps even more compelling are cases of simple inductive or abductive inference, where what you come to know is "manifest" to you, even though you cannot articulate premises of an argument that would license the conclusion as a deductive inference. For example, a person of color might know that what someone said or did was racist, and perhaps be able to point to some features that play a role in licensing the inference, but not be able to fully distinguish such cases from others where the same features are satisfied but the inference is not warranted. Nevertheless, as long as her propensity to believe that people are racist on the basis of these features reflects sensitivity to (enough of) these differences, she can know, on the basis of what was said, that it is racist.

Although we understand the concept of proper practical respect from the paradigm case of knowledge, it plausibly extends exactly to the cases we want, namely those of rational or justified belief whether knowledgeable or not. When we assess a belief that falls short of knowledge as justified, we are condoning the subject's epistemic comportment despite failure. That is, the failure is not something the agent needs to be held responsible for – she was merely unlucky, either because of an unusual environment, or because it was just one of those cases that is compatible with the exercise of a reliable

[19]My thinking on this is indebted to the work of, and personal conversations with, Kurt Sylvan, who argues that rational belief is to be understood in terms of properly valuing the truth. See especially Sylvan (2018). See Miracchi (Forthcoming) for a comparison of our views.

[20]Note that on the knowledge-first approach, this isn't merely a matter of grasping logical connections. For example, there may be cases where grasp of the logical connections between *p*, *If p then q*, and *q* licenses knowledgeable rejection of *p*.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

but fallible competence. If she had failed to manifest proper practical respect for knowledge, then plausibly we should not condone her behavior in this way. We can, then, take cases where the agent manifests proper practical respect for knowledge to be exactly those cases where she exercises her epistemic competence.

As I said, I won't defend this account of rational belief. Instead, I want to show how it can be extended to provide a plausible account of suspension. Note that since suspension isn't a candidate for knowledge, and doesn't constitutively aim at knowledge, it can't be rational or justified *in the very same way* that beliefs are. The account thus rejects *Rational Parity*. However, the agent can still manifest a kind of practical respect for what it takes to know, just a derivative one.

When an agent generally can be characterized as aiming to A (e.g. aiming to know), we can understand withholding from or omitting a performance of A-ing as manifesting a kind of practical respect for what it takes to A. The agent manifests this respect precisely by *not* endeavoring to A. In the case of knowledge, this can be in the form of competently taking yourself not to have (or be able to have) sufficient evidence to settle *p*?, and so intentionally adopting a settled attitude of suspension on *p*?. But it can also be just a matter of having, on the first order, a competence to know. Having good dispositions to withhold is essential to possessing competences to know. (Otherwise, except in very special environments, the reliability condition would fail.) This kind of proper withholding varies from case to case: perhaps what is required is dropping the question altogether, or further engagement on the question by seeking more evidence, thinking harder about the question, and so on. When a person withholds in these ways, we can say that she *demonstrates* her competence to know without exercising it (Miracchi 2015*b*). When a person irrationally suspends, it is because she *should have known* that *p*.

The account thus has the benefit over Sosa's view of explaining the assessability of suspension without over-intellectualizing either belief or suspension. Belief aims at knowledge, not at avoiding error. Suspension need not have an aim at all, precisely because it is in the most basic cases a *non-performance*. Proper suspension on the first order is just a matter of proceeding in accordance with one's first-order competences to know. The account thus satisfies Desideratum 2, that of rationally assessing both first-order and deliberate suspension.

Although it does not explicitly address what broader features of a subject's epistemic comportment are relevant to the assessment of belief and suspension, we have a variety of options for characterizing competences to know that naturally include these broader features. Because epistemic rationality is a matter of practically respecting what it takes to get onto the facts, features outside of what is accessible to us now, or what methods we now have for reliable belief formation, will typically be relevant. The account thus provides a kind of virtue epistemological justification for the inclusion of these other features. The agency involved in properly respecting what it takes to get onto the facts is as broad as agency generally. The account thus resolves the central problem with evidentialism, the inability of the evidentialist to countenance as epistemically relevant features external to how the subject deals with the evidence she happens to have.

It also satisfies Desideratum 3, that the account of suspension be an integrated part of a general account of epistemic rationality. In a broad sense, rational belief and suspension fall into the same category: manifesting proper practical respect for what it takes to get onto the facts. However, it also explains why the rationality of belief is more fundamental and epistemically central. Knowledge is the fundamental aim that structures the epistemic domain, and only in believing can an agent exhibit proper practical respect for knowledge *as the constitutive aim of her performance*. Suspension is epistemically assessable because of how it relates to the aim of achieving knowledge,

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

and indeed is appropriate or not insofar as *not* believing is a way of respecting what it takes to know. This is a benefit over Sosa's view, on which there is total symmetry.

This account has some further interesting features. For example, we should expect plenty of cases where there is *no* epistemically rational option because there is no reason to think that generally, at every moment, we all have the ability to act in a way that practically respects what it takes to get onto the facts. People can get into bad epistemic *ruts*, where they do not have competences to know and are not in environments that can help them develop these competences. For example, someone who grew up in an epistemic environment that offers religious arguments against climate change might not have the resources to rationally believe or suspend because they don't have ways of distinguishing between reliable and unreliable sources.

Often it is supposed to be a drawback of a theory if it has this upshot, presumably because of the implications of some "ought implies can" thesis.[21] However, as discussed in detail above, the kind of ability that is relevant to epistemic assessment is of a broader kind that takes into account both the subject's past in developing certain epistemic habits and their future, including a future where they are subject to epistemic criticism. If ought implies can, it does so only in a much broader sense compatible with this possibility. Sometimes there is no rational course of action an epistemic agent is capable of performing. It is wishful thinking to suppose that in any epistemic encounter there is always *something* we could do that would be beyond epistemic reproach. Typically we have to work hard to be in a position to do the epistemically responsible thing. And because such hard work is within our (diachronic) control, we can be held responsible for believing or suspending in ways that do not properly respect what it takes to get onto the facts, even when we couldn't have done otherwise.

On the other hand, it also allows us to better understand cases of rational disagreement (religious, academic, political, etc.), where neither party to the debate does anything epistemically wrong. These are cases where both are proceeding competently – both are properly respecting what it takes to know, but manifesting such respect requires different commitments – or indeed suspension – for them respectively. This can happen because epistemic competences vary from person to person. Subjects may have different ways of acquiring knowledge; they may also differ in how they gather evidence, what tools they have for reasoning from their evidence, etc. Because our take on the world is always partial (a lesson from standpoint epistemology, cf. Hartsock 1983), and we can sometimes legitimately carve up epistemic tasks differently, there is room for disagreement where neither party is at fault.

For example, it might be rational for an intermediate bird watcher to believe that she sees an American Crow when it is in fact a Fish Crow, but rational for an expert to only believe it is a Fish Crow. The intermediate bird watcher has a lot to learn from understanding why she disagrees with the expert.[22] (This example should make it clear that rational disagreement is not limited to epistemic peers.) Or, consider whether certain counterintuitive consequences of rule utilitarianism warrant its rejection. Deciding this issue involves balancing a number of considerations, and there is plausibly more than one way to competently do so, resulting in differing opinions on the matter.[23] One factor among many here is the trade-off between reliability and

[21]See, e.g. Alston (1988).

[22]Thanks to Shereen Chang for this example.

[23]Interestingly, rational disagreement can then be viewed as a constructive, collaborative effort, that of helping us all get onto the facts by working through the various evidence at our disposal, adjudicating who has expertise on what questions and why, and piecing together a more adequate view of the world. There is no guarantee of convergence, but just the fact that an agent is proceeding competently *before*

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

knowledge possession: epistemic differences can arise when one person is more cautious in forming beliefs than another, even though both are above the threshold of reliability required for competence possession. In such a case, not only may one party permissibly believe while the other permissibly suspends, but a history of systematically doing this may lead to sufficiently different evidence bases, so that what is adequately supported by the more cautious believer is different from what is supported by a less cautious one.

Lastly, it provides a more compelling challenge to Pyrrhonian skepticism, in the spirit of Sosa's challenge. As noted before, the Pyrrhonist is plausibly not *intentionally* withholding belief in *p*. Therefore, we should ask whether her first-order suspension should maintain the kind of default status that *Rational Parity* claims. On my view, this amounts to asking whether the Pyrrhonist is demonstrating proper practical respect for what it takes to get onto the facts in withholding, even on the first order. And the answer here is clearly *No*! She has gone too far. Although we tolerate different degrees of risk aversion, the Pyrrhonist is outside of reasonable bounds for her withholding to count as a demonstration of a competence to know. She is not properly respecting what it takes to get onto the facts because she's not withholding in a way that is sensitive to what it takes to get onto the facts. She takes *nothing* to be sufficient for getting onto the facts.

Because the appropriateness of suspension is not related to avoiding error, but instead to respecting what it takes to know, we can impose a well-motivated limit on what kinds of suspension are rational. To remind the reader: the goal here is not to convince the Pyrrhonian (probably an impossible task) but rather to have a well-motivated theory that does not license the Pyrrhonian position as epistemically superior. We have done this. By our lights, the Pyrrhonian is not playing it safe; she's insensitive to what it actually takes to know. We can and do regularly have knowledge, and her withholding is not a part of that project but a rejection of it.

This explanation of the irrationality of certain kinds of suspension extends to more mundane cases that are of greater importance. The person who suspends on whether there is human-caused climate change, or whether his colleague is behaving inappropriately towards students, is not properly respecting what it takes to know in those domains. Knowledge in the climate change case is easily accessible. In the impropriety case, at the very least settled suspension fails to reflect the ease with which more evidence could put the agent in a position to know. The account thus satisfies Desideratum 1, assessing some cases of suspension as rational and some as irrational, in a way that aligns with both public and academic roles for epistemic assessment.[24]

So, if we adopt my proposal, we can also reject *Rational Parity* in a motivated way, and have a plausible account of suspension that more clearly and directly targets the Pyrrhonist's project. Thus, we can provide a Sosa-like response without requiring that suspending be a performance on par with beliefs. This avoids an implausible commitment about the nature of withholding, one that the Pyrrhonist is unlikely to accept, and it also allows us to do so without attributing a complicated and problematic aim to cases of belief and withholding.

---

an interlocutor brings new considerations to the table does not entail that she may continue to proceed in the same way after these considerations are brought to her attention.

[24]The account also helps us to understand why certain cases of withholding are wishful thinking. Suppose, e.g. Polly wants to go to Morocco, and she aims to know whether she'll be able to afford it. If despite significant evidence that she can't afford the trip she continues to suspend belief, Polly's suspension is irrational. This is because her withholding here is not a matter of respecting what it takes to know in the target domain. Thanks to Baron Reed for suggesting I discuss this example.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

4. Conclusion

I have argued here that evidentialism cannot provide an adequate account of the epistemic assessability of suspension, and correlatively that of belief. The central reason is that our epistemic agency, and so our responsibility, extends much further than to what we do with the evidence we have in the moment. This motivates a virtue epistemological approach, which naturally accommodates such features, over an evidentialist approach, which cannot do so. Then, turning to the question of what kind of virtue epistemological theory we should adopt, I argued for a knowledge-first approach that develops my previous work.

We can now see a connection between the central conclusion of section 2 and the motivations I provided for my approach over Sosa's view. Although I do accept that a *belief* is rational just in case it is an exercise of a competence to know, we need a conception of epistemic rationality that can incorporate other features of our agency as derivatively epistemic. Conceiving of epistemic rationality purely as an exercise of epistemic competence forces us to distort the phenomenon – both to over-intellectualize it and to attribute more similarities to belief and suspension than is warranted. Instead, understanding epistemic rationality as manifesting proper practical respect for what it takes to know plausibly applies in a fundamental sense to exercises of competences to know as such, and also in a broader sense to any relevant features of our agency.

This is a good example of how the study of the rational assessablity of suspension requires us to refine and revise our account of rational belief. A virtue epistemology that treats epistemic rationality solely as a matter of exercise of epistemic competence will lose out on opportunities to take account of these broader features of our agency in epistemic theorizing. The view I have put forward here is a nuanced but still principled account of epistemic agency, one which includes our first-order comportment, including non-performances. It explains suspension as it relates to the aim of achieving knowledge – not avoidance of error – and so commitment to the view that suspension is always the least risky option. It thus assesses some cases of suspension as irrational, providing a compelling response to the Pyrrhonian skeptic.

These upshots I hope are compelling of themselves, but plausibly the more important feature of the account I have offered is that it aligns with a (if not *the*) crucial function of our assessment practices: to hold each other accountable for what we believe and why. Our practices are sensitive to a wide range of agential features, both synchronic and diachronic, retrospective and prospective, and these are plausibly most salient in cases where our epistemic comportment has ethical consequences. The account I have provided here illuminates and legitimizes these features of our practices, not by challenging the distinction between the epistemic and the ethical, but by providing a conception of epistemic rationality broad enough to respect them.[25]

References

Alston W.P. (1988). 'The Deontological Conception of Epistemic Justification.' *Philosophical Perspectives* **2**, 257–99.

Aristotle (1984). *The Complete Works of Aristotle, The Revised Oxford Translation, Vol.* 2. Princeton, NJ: Princeton University Press.

Bump P. (2018). 'Analysis: What We Learned from the Democratic Response to the Nunes Memo and What We Didn't.' *Chicago Tribune*.

Empiricus S. (2000). *Outlines of Scepticism*. Cambridge: Cambridge University Press.

Fleisher W. (2017). 'Virtuous Distinctions.' *Synthese* **194**(8), 2973–3003.

[25]Thanks to Jonathan Jenkins Ichikawa, Baron Reed, Ernest Sosa, and Timothy Williamson.

https://doi.org/10.1017/epi.2019.34 Published online by Cambridge University Press

Goldberg S. (2018). *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford: Oxford University Press.
Hartsock N. (1983). 'The Feminist Standpoint: Developing the Ground for a Specifically Feminist Historical Materialism.' In S. Harding and M.B. Hintikka (eds), *Discovering Reality: Feminist Perspectives on Epistemology, Metaphysics, Methodology, and Philosophy of Science*, pp. 283–310. Dordrecht: D. Reidel Publishing Company.
Littlejohn C. (2012). *Justification and the Truth Connection*. Cambridge: Cambridge University Press.
Lord E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.
Miracchi L. (2015*a*). 'Competence to Know.' *Philosophical Studies* **172**(1), 29–56.
Miracchi L. (2015*b*). 'Knowledge is All you Need.' *Philosophical Issues* **25**(1), 353–78.
Miracchi L. (2017*a*). 'Epistemic Agency and the Generality Problem.' *Philosophical Topics* **45**(1), 107–20.
Miracchi L. (2017*b*). 'Perspectival Externalism is the Antidote to Radical Skepticism.' *Episteme* **14**(3), 363–79.
Miracchi L. (Forthcoming). 'Competent Perspectives and the New Evil Demon Problem.' In F. Dorsch and J. Dutant (eds), *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*. Oxford: Oxford University Press.
NPR (2018). *Fact Check: Read the GOP Memo Released by House Intelligence Committee*. https://www.npr.org/2018/02/02/582828461/fact-check-read-the-gop-memo-released-by-house-intelligence-committee?t=1569676057957.
Peels R. (2017). *Responsible Belief: A Theory in Ethics and Epistemology*. Oxford: Oxford University Press.
Sosa E. (2011). *Knowing Full Well*. Princeton, NJ: Princeton University Press.
Sosa E. (2019). 'Suspension as Spandrel.' *Episteme* in press.
Sosa E. (Forthcoming). 'Suspension and Evidentialism.' *Synthese*.
Sylvan K. (2018). 'Veritism Unswamped.' *Mind* **506**, 381–435.
Toole B. (2018). *Knowledge and Social Identity*. PhD thesis, University of Texas at Austin.
Williamson T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
Williamson T. (Forthcoming). 'Justifications, Excuses, and Skeptical Scenarios.' In F. Dorsch and J. Dutant (eds), *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*. Oxford: Oxford University Press.

Lisa Miracchi is an Assistant Professor of Philosophy at the University of Pennsylvania. Her research concerns issues in philosophy of mind, cognitive science, artificial intelligence, and epistemology.

Cite this article: Miracchi L (2019). When Evidence Isn't Enough: Suspension, Evidentialism, and Knowledge-first Virtue Epistemology. *Episteme* **16**, 413–437. https://doi.org/10.1017/epi.2019.34