Philosophical Studies
https://doi.org/10.1007/s11098-024-02247-y

What is reasonable doubt? For philosophical studies special issue on Sosa's 'epistemic explanations'

Lilith Mace[1] · **Mona Simion**[1]

Accepted: 9 October 2024
© The Author(s) 2024

Abstract
This paper develops and defends novel accounts of accurate and reasonable doubt. We take a cue from Sosa's telic epistemic normative picture to argue that one's degree of doubt that p is accurate just in case it matches the level of veritic risk involved in believing that p. In turn, on this account, reasonable doubt is doubt that is generated by a properly functioning cognitive capacity with the function of encoding veritic risk.

Keywords Doubt · Suspension · Proper function · Cognitive capacity · Knowledge-first

1 Introduction

When our epistemic life is nice and simple, we believe things that we know to be true, and we suspend belief otherwise. We are cognitively limited creatures, however, so our epistemic life doesn't always pan out like this. Epistemic support comes and goes, sometimes evidence leads us astray, defeaters come in right when we thought things were settled, we are at high risk of error although we get it right, and many times we doubt our conclusions and need to go inquire further to figure things out. Luckily, however, we are complicated cognitive machines, that can cope with this mess: We believe, we judge, we form credences and doubts, we suspend, we

✉ Mona Simion
mona.simion@glasgow.ac.uk

Lilith Mace
lilith.mace@glasgow.ac.uk

[1] Cogito Epistemology Research Centre, University of Glasgow, 69 Oafield Avenue, Glasgow G12 8LP, UK

Published online: 08 November 2024

reason and weigh evidence and defeat. This paper is concerned with the epistemic propriety of one aspect of this mess: the epistemic propriety of doubt.

Until recently little has been said about the justification of non-affirming attitudes, such as suspension and doubt.[1] In recent work, Ernie Sosa takes up this task: Sosa develops a telic virtue-theoretic account of epistemically permissible suspension, and ventures to thereby explain the epistemic impermissibility of unwarranted sceptical doubts. On Sosa's view, in a nutshell, one is permissibly suspended on p just in case one suspends on p based sufficiently on one's lack of the competence required in order to answer that question whether p aptly. In turn, according to Sosa, what goes wrong in cases of sceptical doubt is that the sceptic could aptly judge on many matters, but she fails to do so.

We share Sosa's telic normative commitments, and we think the picture defended in *Explanations* is very illuminating; however, we worry it is vulnerable to objections from cases of unreasonable doubt sourced in lack of competence. What we plan to do in this paper is two-fold: first, we outline our worry for Sosa's picture and look at a few possible moves that could be made on Sosa's behalf. Second, we venture to develop our own telic view of epistemically permissible doubt. Our account unpacks telicity as etiological functionality, and epistemic norms as functional norms. That being said, since the functionality at stake is construed as pertaining to our cognitive capacities, we think Sosa can usefully and easily incorporate our proposal (or a close cousin thereof) in his framework.

2 Sosa on suspension, recklessness, and scepticism

Sosa's virtue epistemology is a normative framework for the evaluation of attempts (henceforth also 'telic normativity'). Attempts have constitutive aims. As a result, we can ask whether or not a given attempt is successful. We can also ask whether a given attempt is competent, i.e. produced by an ability to attain the attempt's aim. Finally, we can ask whether a given attempt is apt, i.e. successful because competent.

Virtue epistemologists[2] standardly take beliefs to be attempts that have truth and/or knowledge as their constitutive aims. On Sosa's truth-based picture, we can ask whether beliefs are successful, i.e. whether they are true. In addition, we can also ask whether they are competent, i.e. whether they are produced by an ability to believe truly and whether they are apt, i.e. true because competent.

According to Sosa, the above gives us the basic account for first order evaluations of attempts. Crucially, however, Sosa does not take this to be the whole story. Rather, he countenances two further types of aptness, alongside first order aptness, or 'animal' aptness as Sosa calls it. These additional types of aptness are 'reflective' and 'full' aptness. Attaining these further types of aptness requires accurate

[1] For exceptions see e.g. Mace Forthcoming, Friedman (2017), Miracchi (2017), Lord and Sylvan (2021), Simion (2024), Williamson (2005).

[2] For truth-based versions of the view, see e.g. Sosa (2021), Greco (2010) and Pritchard (2012). For knowledge-based accounts see e.g. Kelp (2017), Miracchi (2015).

and indeed apt attempt at a higher order, in addition to animal aptness. In a nutshell, the thought is that attempts will rise to these higher levels of aptness only if, alongside animal aptness, one has aptly ascertained that one's attempt is free from any relevant risk one may be running: one must have arrived at an apt awareness that one's attempt would be apt. While animal aptness in conjunction with apt risk assessment will be enough for reflective aptness, full aptness additionally requires that first and second order aptness are connected in the right way: one must be guided to animal aptness by one's reflectively apt risk assessment.

It comes to light that there are a number of normative properties that attempts can enjoy. Crucially, according to Sosa, full aptness enjoys special status among these properties. More specifically, according to Sosa full aptness is the fully desirable status for attempts and that attempts fall short unless they attain full aptness. Moreover, he is also clear that this claim holds with full generality: Any attempt attains fully desirable status qua attempt if and only if it is fully apt; and it falls short qua attempt if and only if it isn't.

According to Sosa, various psychological categories—most importantly, belief and judgment—are species of affirmation and, as a result, attempts. (Sosa's main interest is with affirmations with a specifically epistemic aim, which, at a minimum, involves truth). While Sosa countenances a variety of psychological categories with epistemic aims, his main focus is on judgment (and judgmental belief). Judgment differs from other psychological categories in that it has a particularly robust epistemic aim: judgment aims not only at truth, but at aptness.

To understand this normative requirement on judgment, Sosa asks us to consider Diana, the huntress: As Diana surveys a landscape in search of game, she may see prey in the distance (in good light and calm wind). If a shot is too risky, it is ill-advised. A shot, then can attain quality in being well, rather than negligently selected. An aiming, then, is assessable by reference to how likely it is to succeed (relative to one's possession of the pertinent competence), so as to avoid recklessness, and also assessable by reference to how negligent (or not) it may be.

Similarly, according to Sosa, for a judgment to be apt, more is required than merely apt affirmation. What is needed for apt judgement is that one is guided to aptness by apt risk assessment. An apt judgement is a fully apt affirmation. Consider Diana again. Diana's archery shots can be more or less well selected. When she spots some prey, Diana can properly aim as follows: to make an attempt on that target if and only if the attempt would succeed aptly. Accordingly, there are two ways in which Diana can fall short, with regard to this aim: she could make an attempt on the target when she would not succeed aptly—because, maybe, the shot would be too risky, given the wind. But she could also fail in her attempt by failing to make an attempt (on the target) when one would succeed aptly.

So, in a nutshell, according to Sosa, forbearing from attempting can itself be an attempt with an aim: that of attempting if and only if the attempt would succeed aptly. This is the place of forbearing in telic normativity.[3]

[3] A similar idea is developed in Broncano-Berrocal (2016, 2018). According to Brioncano-Berrocal, we host precautionary abilities, that are not aimed at completing tasks, but play a precautionary or protective role instead. They are dispositions to stop the completion of a task by a task-completion ability when the

Sosa thinks that epistemic forbearing is what constitutes deliberative suspension of judgment, which is an attempt in its own right, one that shares with judging an epistemically distinctive aim: the aim of affirming alethically (positively or negatively) if and only if that affirming would be apt (and otherwise suspend). Just like Diana, a subject can fall short of epistemically apt forbearing in one of two ways: she could judge when her judging would be inapt—because too risky. This is the case of Barney in the Fake Barn County. But she could also fail to judge when she should, because she would thereby judge aptly. Here is Sosa on unreasonable sceptical doubt:

> […] the propriety of suspension turns on the agent's proper assessment of the epistemic risks. As an agent you have an aim in suspending. How appropriately you suspend is hence *pro tanto* directly proportional to how competently you assess the pertinent risk. And I can see no reason to suppose a priori that *taking the risk to be too high* should be preferable by default over *taking the risk to be appropriate* (2021, Ch. 3).

On Sosa's view, one properly suspends belief on a question if and only if one suspends based sufficiently on one's lack of the competence required in order to answer that question aptly (2021, 85). Thus, suspension is a kind of "spandrel": an "inevitable byproduct" of aiming for apt success; in this case, for apt alethic affirmation (ibid.). In order to affirm alethically when, and only when, such affirmation would be apt, it is necessary that one suspends judgment when such judgment would not be apt.

In this, Sosa's account promises to have the resources needed to account for both reckless and overly-cautious judging. Indeed, Sosa takes the view to explain what goes wrong in the case of the Pyrronist sceptic. He writes.

> This bears on Pyrrhonian skepticism. Why should it be thought that the stance of the Pyrrhonist must be granted any sort of default status? Why should we demand from the dogmatist any special effort to overcome that default status, while allowing the skeptic to sustain his position until the dogmatist manages to displace him? (2021, Ch. 3)

The Pyrrhonist, the thought goes, errs in the direction of overcautiousness: she could aptly judge on many matters, but she fails to do so.

3 The problem of unreasonable doubts

We have seen that Sosa's telic epistemological picture promises to have the resources to explain both judgment that is epistemically defective because too reckless, and what goes epistemically wrong with overly cautious forbearing from

Footnote 3 (continued)
conditions are not suitable for it (e.g. a world-class archer might be disposed to refrain from delivering arrows whenever the wind speed is too high to hit the mark).

judging in low-risk environments. In what follows, we will express some worries on the latter front; that is, we will argue that Sosa's picture struggles to accommodate an important class of cases of unreasonable doubt: cases in which doubts are intuitively epistemically problematic even though the doubting subjects are not in a position to aptly judge on the issue.

To begin with, consider the following classic case of unreasonable doubt from Barber & Gordon:

Unreasonable Juror (UJ): … [I]n the jury room, one man reacted to the concept of 'reasonable doubt' as if it were a challenge to his ingenuity. It meant, he insisted, that we were to see if we could think of any possible alternative explanation of events, and he could—somebody had 'planted' the garment in the girl's bag. It was pointed out to him that even the defence had not put forward this explanation. This didn't matter, he said, perhaps they hadn't thought of it. There *must* be reasonable doubt if you could construct another theory, after all, it wasn't physically impossible, was it? It was now pointed out to him that although it wasn't physically impossible, his explanation was not based on a single scrap of evidence. Who did he think had done the 'planting', the store detective? 'A person or persons unknown,' said the odd man out, proudly. (1976, 76).

The Unreasonable Juror doubts that the defendant committed the crime insofar as he can think of any (nomically) possible alternative explanation. We take it that the correct diagnosis of this case is that the UJ is epistemically in the wrong: UJ doubts that the defendant committed the crime, even though he shouldn't: he has plenty of evidence to believe that she is guilty. Does Sosa's framework have the resources needed to accommodate this result?

At first glance, Sosa's account of suspension seems to be exactly the kind of account that serves to explain what is going wrong in this case: after all, Sosa could argue, UJ is impermissibly suspended by the lights of telic epistemic normativity: he is in a position to competently affirm that the defendant committed the crime, but he fails to do so: risk assessment need not go as far in the space of possibility as UJ does. After all, risk is a matter of distance in the space of possibility (Pritchard, 2015): an event is only risky if it occurs in a world that's close enough to our own. Just like Diana need not take the risk of e.g. a meteor hitting Planet Earth into account when assessing her likelihood to shoot aptly, neither should UJ go through all of these far-fetched scenarios before believing the defendant committed the crime. UJ is, like the Pyrrhonist sceptic, in epistemically impermissible forbearing from believing, thereby.

On second thought, however, there will be some difficulties with this reply:

First, note that, even if it is plausible to construe the case as one where UJ suspends on whether the defendant committed the crime, that is not all there is to it. UJ is not merely suspending on the issue, he's also actively engaging in a negatively-valenced attitude to the issue: he is doubting that the defendant did it. As such, any correct account of what is going on in this case requires resources to explain both the impermissibly of suspending, and the impermissibly of the associated instance of doubt. A telic-normative account of epistemically permissible

doubt is required. Since Sosa hasn't offered one, his picture is bound to not give us everything that we need here.

On the other hand, however, one might think this is also good news for Sosa: after all, if his account does successfully explain what is wrong with UJ's suspension, it is up to virtue epistemology scholarship to develop a corresponding account of permissible doubt in order to fully account for the epistemic negative phenomena present in this case.

The problem with this response, however, is that it remains too weak to explain the data. To see why, consider a variation on the Unreasonable Juror case, UJ*: say that the underlying reason why UJ* engages in such unreasonable sceptical doubts has to do with his (secret!) deep seated bias in favour the defendant: say that UJ* is a racist and a sexist, and the defendant is a white man just like UJ*, while the victim is a black woman. Or, to remove any suspicion of intuitions getting confused between moral and epistemic normativity, take UJ*'s bias to be a cold rather than hot bias (for instance, take UJ* to favour the defendant in virtue of confirmation bias): nothing hinges on this.

Note that in this variation of the case, UJ*'s withholding meets the sufficiency direction of Sosa's (envisaged) account of permissible withholding: after all, UJ* could not, in virtue of his bias, form a competent judgment on whether the defendant committed the crime or not. Since he is thereby suspending, Sosa's view predicts he is epistemically permissibly suspending. Once more, the account fails to accommodate the intuition that something epistemically problematic is happening in this case.

Finally, note that we can also simply stipulate that UJ* is not suspended in the case above, but rather lacks any belief on the issue, and merely hosts a negatively valanced doxastic attitude towards the proposition that the defendant did it: he doubts it. If suspension is not present, neither is impermissible suspension. As such, the question that arises is: what resources in Sosa's telic normativity are there still available to explain what goes wrong in the new UJ* case? Note that it will be harder in this case to push the issue on the shoulders of the doubt theorist: after all, intuitively, independently of the doubts he is harbouring, this is a clear case of impermissible belief withholding; a case in which UJ should believe that the defendant did it, but fails to do so. Can Sosa accommodate this result in his framework for assessing the normativity of belief and judgment? At first glance, it would seem that, for all we have been told so far, this will be difficult. After all, on Sosa's view, the normativity of belief is the normativity of attempts. However, the UJ* case is a paradigmatic case in which no attempt is being made even though an attempt should have been made: UJ* has all the epistemic support he needs in order to aptly affirm that the defendant committed the crime, but fails to do so. Since no attempt at apt judgment is made, however, an epistemology build on the normativity of attempts risks remaining silent on what is going wrong in this case.

Not so fast, though: After all, Sosa may argue, while forming a suspension is one paradigmatic case of forbearing from affirming, withholding belief is yet another variety thereof. Modelling epistemically permissible belief withholding on the normativity of permissible forbearing, then, we get an account on which withholding belief is epistemically permissible if and only if one withholds based sufficiently on one's lack of the competence required in order to answer that question aptly (2021,

85). It is easy to see, however, that this response will lead us straight back into the trouble of the biased unreasonable juror that we identified for Sosa's account of suspension: after all, UJ could not, in virtue of his bias, form a competent judgment on whether the defendant committed the crime or not. Since he is thereby withholding, Sosa's view predicts he is epistemically permissibly withholding. But he is not.

4 Epistemic risk and reasonable doubt

In the previous section we looked at a case of unreasonable doubt and argued that Sosa's telic epistemological framework struggles to explain the corresponding epistemic impermissibility intuition. In what follows, we will develop our own proposal for a telic account of reasonable doubt. The account is functionalist, in that it takes the function of doubt in our cognitive system to explain its epistemic normativity. Here it goes:

On the picture we favour, generating knowledge is the epistemic etiological function of our cognitive system.[4] The function is etiological in that our cognitive system has generated knowledge in the past in our ancestors, which was beneficial, and which contributes to the explanation as to why cognitive systems continue to be replicated in creatures like us (Millikan, 1984; Graham 2000; Simion, 2018, 2019a, 2019b, 2019c). In turn, cognitive processes—such as forming beliefs, judgments, suspensions, credences, doubts, assertions, or pieces of reasoning—are aimed—either directly or indirectly—at fulfilling this knowledge-generating function. The difference between direct and indirect conduciveness to function fulfilment lies with achievability: since e.g. beliefs and judgments can be knowledgeable, in a way in which things like e.g. credences, suspensions, and doubts cannot, belief and judgment formation are aimed directly at fulfilling the function of our cognitive system (generating knowledge); while forming credences, doubts, and suspensions is aimed at knowledge indirectly: they are attitudes held in the service of knowledge, that support our cognitive system in generating knowledge via serving subsidiary epistemic functions, but that are not in the running for knowledge (Simion, 2024).

The function of our cognitive system is to produce knowledge. Doubt has a subsidiary function, which supports the knowledge-producing function of the system. This function is to act as what Williamson calls our "cognitive immune system", which "protect[s] our conception of the world from harmful errors" (2005: 681).

Crucially, in order for doubt to fulfil its protective function, it is not enough that it merely prevent the uptake of beliefs that would actually be false. This is for two reasons. First, a cognitive immune system that is weak enough to let in true beliefs that could easily have been false will inevitably be weak enough to let in myriad false

[4] One of us has defended this claim extensively in previous work (e.g. Kelp & Simion (2017, 2020, 2023), Simion (2016, 2018, 2019a, 2019b, 2019c, 2021a,2021b, 2021c, 2021d, 2023 2024), Simion and Kelp (2018a,2018b)), rehearsing these arguments falls outside the scope of this paper. See also Williamson 2000. For etiological theories of function, see e.g. Milikan (1984), Neander (1991), Godfrey-Smith (1994).

beliefs, too. Second, true beliefs that could easily have been false are not candidates for knowledge.

Just as the body's immunity system can be overactive, attacking healthy tissue alongside genuine threats, so too can our cognitive immunity system be overactive. Instead of protecting us from forming or holding beliefs that are at genuine risk of error, overactive doubts can prevent us from forming beliefs that are well supported by our evidence and thereby candidates for knowledge. This is the predicament of the sceptic, who doubts that she has hands, because she takes seriously the possibility that she is a brain in a vat.

Consider, in contrast, Barney, the Fake Barn County visitor: were Barney to doubt that what he's looking at is a barn, his doubt *would* track a genuine risk of error. His doubt, in this, would be an accurate doubt. A belief, in Barney's case, that what he's looking at is a barn, is a belief that is at genuine risk of being false: fixing the method by which Barney arrives at his belief, he would form a false belief in some close world.

We will follow Duncan Pritchard (2015, 2016)[5] in thinking that what makes an event risky is that, while it fails to obtain in the actual world, there are—keeping the initial conditions for that event fixed—close possible worlds in which this event obtains. Further, we take the level of risk involved in an event to be determined by how far is the closest world at which the risk event obtains (within this space of close possible worlds); that is, by how modally close the target risk event is in the space of close possible worlds.[6] Roughly, where the risk event is modally very close, then it is high-risk, but where it is not modally very close, then it is low-risk. (This way to put it is rough because one might think that the number of close worlds at which the event obtains also matters for how high the risk is—but this complication will not matter for us here.)

Crucially, in making these assessments, we need to keep relevant initial conditions fixed across cases. For instance, when assessing the risk of a plane crash, we are interested in those close possible worlds in which one continues to catch a plane. When it comes to epistemic risk, this will translate as follows: keeping the method of belief formation fixed, when the belief is false in a close world, what we have is veritic epistemic risk (risk that the belief in question is false). The closer is the closest world in which the (would-be) believed proposition is false, the higher the level of veritic risk: When it is false at a very close world, we get high veritic epistemic risk. If, in contrast, the risk event doesn't occur at a very close world, then it is low risk. Finally, where the risk event—i.e. the false belief—does not occur in any modally close world, then it is not risky.

Risk comes in lower and higher degrees. Correspondingly, so does doubt: since we're in Scotland, we doubt that it's going to be sunny and dry tomorrow more than we doubt that the Tories will win the elections. That is because the risk

[5] Competing accounts of the nature of risk include the probabilistic account (Hansson 2023) and the normic account (Ebert, Smith and Durbach 2020). While we work with the modal account here, not much hinges on this. The reader is free to replace with their favourite account of risk.

[6] For more on the notion of an easy possibility, see e.g. Sainsbury (1997).

involved in believing that it's going to be sunny and dry tomorrow in Scotland is higher than the risk involved in believing that the Tories will win the elections. Compatibly, we doubt both that it's going to be sunny and dry tomorrow and that the Tories will win the elections: both of these beliefs would incur a pretty serious level of veritic risk.

This suggests the following account of accurate doubt:

> **Accurate Doubt**: One's degree of doubt whether p is accurate if and only if it matches the level of veritic risk involved in believing that p.

This view allows us to explain the fact that doubts can be more or less accurate: our doubting that it will be dry and sunny tomorrow is more accurate than the sceptic's doubt concerning the existence of the external world. That is because our doubt tracks the degree of closeness of a world in which it rains better than the sceptic's doubt tracks the closeness of the world where we're brains in vats. We can thereby rank doubts with regard to their degrees of accuracy, and compare them accordingly, by looking at how well they track veritic risk: the closer a degree of doubt that p tracks the level of veritic risk involved in believing p, the more accurate the doubt is.

The fact that doubt admits of comparisons suggests that 'doubts that' is a gradable predicate. Gradables, in turn, come in two varieties: absolute gradables (such as 'is closed' or 'is impure') and relative gradables (paradigmatic examples include 'is tall' and 'is flat' (Kennedy 2007)). Absolute gradables differ from relative gradables in what makes for a proper outright attribution: in the case of relative gradables, proper attribution is context sensitive: men who are 6 feet tall are outright tall in Italy but not in Sweden. In contrast, absolute gradables' attributions are not context-sensitive. Rather, their proper attribution depends on their arguments displaying either a maximum or a minimum degree of the property they describe: a door is only closed if it is fully closed, and gold is impure if it is at least minimally impure.

It is easy to see that 'doubts that' pertains to the latter category of absolute gradables: it is enough that I doubt that p to a very small degree for it to be true that I doubt that p simpliciter. In this, 'doubts that' works like 'is impure'. Compatibly, however, pragmatic contextual factors might make assertability of outright doubt claims problematic. Take a simple relevance implicature, for instance: it seems problematic for a doctor to tell their patient that he doubts the safety of vaccines in a case in which he harbours a very small degree of doubt. That is not because it isn't strictly speaking true that he doubts it, but rather because asserting it outright at this context would trigger the false implicature that he harbours a degree of doubt that is high enough to be relevant at this context.

Moving on: We have said that, if Barney were to have a high level of doubt that there's a barn in front of him when visiting Fake Barn County, his doubt would be accurate, in that it would successfully track the high veritic risk present at that context. Note, however, that, as the story goes, should Barney actually harbour said accurate doubt, that would be quite strange—epistemically. After all, Barney has no reason to believe there's any risk involved in believing that there's a barn in front of him, since he has no idea that he finds himself in Fake Barn County to begin with. Should he highly doubt that there's a barn in front of him, then, accurate as this doubt may luckily be, it would not be a reasonable doubt for Barney to harbour.

His doubt would fail to successfully encode the high veritic risk at this case: rather, it would be merely luckily matching it. Our cognitive capacities, however, are not merely aimed at matching the world: they are aimed at encoding it. Norman the clairvoyant is a reliable world matcher, but his beliefs are epistemically defective nevertheless.

For just this reason, the accuracy of a doubt can come apart from whether that doubt is reasonable, or justified.[7] The reasonableness of a doubt, we submit, will depend on one's evidential situation with respect to what the risk facts look like, and on how well the process of forming the relevant doubt through update on this evidence went. In what follows, we unpack this claim.

We said that our cognitive capacities have etiological epistemic functions. Now, one important feature of etiological functions is that they have normative import. For instance, once the heart acquires the function of pumping blood, there is such a thing as a properly functioning heart—i.e. a heart that works in a way that, in normal environmental conditions, reliably enough pumps blood—and such a thing as a malfunctioning heart—a heart that fails to work in this way. In the case of the human heart, normal functioning involves beating at a certain rate (roughly, between 60 and 100 beats per minute when resting). The conditions of proper functioning are thereby normatively governing hearts: hearts ought to work in a way that, in normal environmental conditions, reliably leads to function fulfilment—i.e., to pumping blood: they ought to beat at a rate of 60–100 beats per minute when resting (Graham 2012; Simion, 2019a, 2019b, 2019c).

With these points in play, let's move on to what it takes for a doubt to be reasonable, or justified. The key idea we would like to start with is that our doubt-generating cognitive capacities have the function of encoding veritic risk. This function enables our doubts to act as our "cognitive immunity system", and protect us from forming or holding beliefs that are at genuine risk of error, and thereby not in the market for being knowledgeable.

That our doubt-forming cognitive capacities have the function of encoding veritic risk is independently plausible, but also predicted by the etiological requirement on function acquisition: it is plausible that our cognitive capacities have generated doubts that encoded veritic risk in the past, which was beneficial epistemically (for protecting our cognitive system from forming beliefs that are at genuine risk of error), which, in turn, contributes to the explanation of the continuous existence of said capacities.

Once more, importantly, *encoding* veritic risk doesn't simply amount to matching veritic risk. In the case of belief, you may arrive at a belief that is true as a result of the operation of two biases, which luckily cancel out each other's effect. In the same way, in the case of doubt, you can arrive at a doubt that matches your veritic risk as a result of the operation of two biases, which luckily cancel each other out. In the

[7] For the purposes of this paper, we use 'reasonable' and 'justified' interchangeably. Nothing hinges on this. Our preferred terminology is the terminology of justification. However, the traditional terminology employed for talking about justified doubts within and outside of philosophy is reasonableness terminology.

case of belief, your belief forming capacity will not have thereby fulfilled its function: your belief matches the world, but it falls short of knowledge. Similarly, in the case of doubt, your doubt forming cognitive capacity will not have thereby fulfilled its function: your doubt matches veritic risk, but it does not encode it.

Now, recall that functions have normative import in that they give rise to norms of proper functioning. In the case of the cognitive capacities under consideration, this means that these norms tell us what cognitive capacities ought to do. The result that we get is that the cognitive capacities involved in the formation of doubts ought to function properly, that is, they ought to function in the way that, in normal environmental conditions, reliably enough leads to function fulfilment—i.e., to encoding veritic risk.

Finally, our key claim is that the reasonableness, or justification of a doubt corresponds to the satisfaction of this epistemic norm of proper function. In this way, we get the following account of reasonable doubt:

Reasonable Doubt: One's doubt that p is *prima facie* reasonable iff it is formed via a properly functioning cognitive capacity that has the function of encoding the veritic risk involved in believing p.

Crucially, the proper functioning of the relevant cognitive capacities will imply proper update on available evidence of risk. The proper function of our cognitive system in general involves the proper evidential uptake and update: After all, that's how our cognitive system works when it produces knowledge in normal conditions. If you are looking at a blue dot on a white background and you do not take up the information that there is a blue dot in front of you, your cognitive system is not functioning properly. This suggests that the proper functioning of our cognitive system is input-dependent, in that it implies proper trigger response. One illuminating analogy here is the proper functioning of the lungs. One important function of lungs is to supply their host organisms with oxygen. Crucially, this involves not only the transferral of oxygen into the bloodstream, but also the taking up of oxygen under certain trigger conditions. Accordingly, there are two ways in which lungs can be malfunctioning: by deviating from the normal way of transferring oxygen into the bloodstream and by deviating from the normal way of taking up oxygen. That is to say, one way in which lungs can malfunction is by failing to take in oxygen in the presence of the right trigger conditions.

If one's body of available evidence suggests that the risk facts are thus-and-so, then one's doubt is reasonable to the extent that one's properly functioning cognitive capacities with the function of encoding veritic risk pick up the relevant evidence and generate a corresponding degree of doubt. Reasonableness, then, like accuracy, has to do with a similarity ordering on worlds. But while accuracy centres on the actual world, reasonableness centres on the set of worlds that one's evidence suggests are close. S's doubt, then, will be a reasonable doubt insofar as it is formed via a properly functioning cognitive capacity that has the function of encoding the veritic risk involved in believing p, where proper function implies uptake and update of easily available evidence of degrees of veritic risk.

The view has a number of theoretical advantages, including: that it is naturalistically respectable—in that it unpacks the normativity of doubts in etiological

function-theoretic terms; that it allows us to accommodate the presence of reasonable doubt in scenarios in which no actual veritic risk obtains, but where the subject has misleading evidence that it does; conversely, that clairvoyant doubters—i.e. cognizers whose doubts reliably match veritic risk by luck—are not reasonable doubters: uninformed doubting Barney is accurately but not reasonably doubting. More down-to-earth, the view also explains, for instance, why subjects doubting well-established scientific facts—such as the occurrence of anthropogenic climate change, or the safety of vaccines—are often, but not always, unreasonable: most of us inhabit epistemic environments where our evidence suggests that believing these well-established scientific facts does not come at any veritic risk. However, for others among us, inhabiting polluted epistemic environments, misleading evidence suggests that the world at which these claims are false is close.

This account of doubt can also be put to good practical work. For example, legal theorists and practitioners have had difficulty, historically and up to the present day, in explicating the 'beyond reasonable doubt' standard of proof, both among themselves and to jurors (see Lauden, 2011 for discussion). A number of philosophers have suggested that the standard should therefore be replaced, and have offered alternatives that make use of philosophical notions supposedly better understood: for example, Thomson (1986) proposes a causal standard of proof, whereby jurors should find for p if and only if their judgement that p is caused by the fact that p; Enoch et al. (2012) suggest the standard is replaced by a sensitivity standard, whereby jurors find for p if and only if their judgement that p is sensitive to p's truth, in that they wouldn't judge that p if p were false; and Pritchard (2018) argues for replacing the standard with a safety standard, whereby jurors find that p if and only if their judgement that p is safe, i.e. could not easily be false. Yet our account of reasonable doubt can help shed light on the existing 'beyond reasonable doubt' standard of proof, adding clarity to legal theory without advocating for major structural change in existing legal practice.

One issue surrounding the interpretation of the 'beyond reasonable doubt' standard is that jurors misunderstand the restriction on veritic risk-possibilities that reasonable doubts are supposed to encode. Recall UJ: "[..] he insisted, that we were to see if we could think of any possible alternative explanation of events, and he could—somebody had 'planted' the garment in the girl's bag. It was pointed out to him that even the defence had not put forward this explanation. This didn't matter, he said, perhaps they hadn't thought of it. There *must* be reasonable doubt if you could construct another theory, after all, it wasn't physically impossible, was it?" (1976: 76). One straightforward way to understand the mistake of the unreasonable juror is to say that he mistakenly takes relevant risk to extend throughout the space of nomic possibility, and reasonable doubt to follow suit. After all, on this juror's view, S's doubt that p is reasonable if and only if p is false in some nomically possible world. But on our view, this is not the case. Not all propositions that are false in some nomically possible world are propositions that would be reasonable to doubt. There will be many propositions such that, although it is nomically possible that they are false, they are true in all worlds that the relevant body of evidence are close. Relative to the jury's body of evidence—which presumably includes propositions that speak in favour of the defendant having placed the garment in her bag

herself—it could not easily have happened that someone else placed the garment there. As such, the juror's doubt is unreasonable.

5 Practical risk and reasonable doubt

Given that our account of the normativity of doubt appeals crucially to epistemic risk, the reader may now wonder whether, and to what extent, practical considerations encroach on the normativity of doubt. Where the practical consequences of false belief would be severe, surely the risk is higher than it would be if the stakes were lower; if the risk is higher, certain doubts that would otherwise have been unreasonable may count as reasonable. And this might seem intuitively correct: it is natural to think that whether a doubt about the defendant's guilt is reasonable depends on what would happen to her, were she convicted. If she would receive a fine, the possibility that the garment was planted in her bag cannot underpin a reasonable doubt; if she would be executed, it can.

Furthermore, if this is correct, it would seem as though Sosa's framework may have the resources needed to accommodate cases like UJ and UJ* within what he calls 'intellectual ethics' (2021)—the domain that, while external to the epistemic domain proper, deals with moral and practical norms that affect inquiring and evidence gathering.

We want to resist this intuition. We are giving an account of what makes doubt epistemically reasonable. As such, the relevant risk is epistemic risk: the risk of believing in error. How epistemically reasonable a doubt is, is entirely a matter of how close the error possibility that the doubt encodes is, relative to one's body of evidence: if one's evidence suggests that the possibility is close, then a doubt that encodes that possibility is reasonable; if one's evidence suggests that the possibility is not close, the doubt is unreasonable. That false belief on a particular occasion would have bad practical consequences does not affect this ordering on worlds, and so makes no difference to the level of epistemic risk in play. As such, practical considerations do not make a difference to the epistemic reasonableness of doubts.

However note that whether an epistemic risk is high or low is orthogonal to whether that risk is (practically, morally, or all-things-considered) acceptable. This is true of risk in general: sometimes high risks are acceptable, and sometimes low risks are unacceptable. The risk that one's body will reject an organ transplant may be high but acceptable, if the alternative to receiving the transplant is death; while the risk that an oil rig will explode might be incredibly low but unacceptable, given the devastation that would follow. Similarly for epistemic risk: a fairly close possibility of error may be acceptable for trivial matters on which nothing practical depends, such that a subject who entertains doubts on the matter would strike us as pedantic.

For example, a subject who ruminates on the possibility that she is misremembering that Phil Mitchell is the older brother to Grant, rather than vice versa, may be entertaining a reasonable doubt: perhaps she knows that has a bad memory, so could easily be misremembering. Yet her doubt is of so little practical import, and concerns a matter so trivial, that we think she ought not entertain it: the level of

epistemic risk, although high, is acceptable. In contrast, very low epistemic risk can be unacceptable, where the stakes are high enough. Suppose that I know that the code to defuse a bomb is the product of 25 and 90, which I mentally calculate, correctly, to be 2250. Further, I know that I am good at mental mathematics, so my evidence suggests that I couldn't easily be wrong in this case. To doubt that the code to defuse the bomb is 2250 would then be unreasonable: my evidence suggests that this possibility is not close. Yet given the practical costs of false belief—the bomb will go off!—it may be that I ought to entertain this doubt, and double-check my answer with a calculator before inputting the code.

As should be clear, whether a level of epistemic risk (or indeed any kind of risk) is acceptable is an all-things-considered matter. As such, practical and moral considerations can make a difference to whether an epistemic risk is acceptable. It may sometimes be the case that a very low level of epistemic risk, and indeed a low level of epistemic risk relative to some body of evidence, is all-things-considered unacceptable. Insofar as entertaining a doubt in such cases motivates behaviour that reduces the level of epistemic risk in play to an all-things-considered acceptable level, entertaining the doubt may be the all-things-considered best thing to do. But that does not make the doubt *epistemically* reasonable.

6 Conclusion

We have proposed and defended a telic account of reasonable doubt. The view is functionalist, in that it takes the normativity of doubt to be borne out by the etiological function of our doubt-generating cognitive capacities. On this account, one's doubt is reasonable just in case it is generated by a properly functioning cognitive capacity with the function of encoding veritic risk. We have also argued that Sosa's telic normative framework could benefit from incorporating our account of reasonable doubt.

Funding Funding was provided by HORIZON EUROPE European Research Council (Grant No. 948356, KnowledgeLab).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Barber, D., & Gordon, G. (1976). *Members of the jury*. Wildwood House.
Broncano-Berrocal, F. (2016). A robust enough virtue epistemology. *Synthese, 194*, 2147–2174.

Broncano-Berrocal, F. (2018). Purifying impure virtue epistemology. *Philosophical Studies, 175*(2), 385–410.
Ebert, P., Smith, M., & Durbach, I. (2020). Varieties of risk. *Philosophy and Phenomenological Research, 101*(2), 432–455.
Enoch, D., Fisher, T., & Spectre, L. (2012). Statistical evidence, sensitivity and the legal value of knowledge. *Philosophy and Public Affairs, 40*, 197–224.
Friedman, J. (2017). Why suspend judging? *Noûs, 51*(2), 302–326.
Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs, 28*, 344–362.
Graham, P. (2012). Epistemic entitlement. *Nous, 46*(3), 449–482.
Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
Hansson, S. O. (2023). Risk. In E. N. Zalta, & U. Nodelman (Eds.), *The Stanford encyclopaedia of philosophy (summer 2023 edition)*. https://plato.stanford.edu/archives/sum2023/entries/risk/
Kelp, C. (2017). Knowledge-first virtue epistemology. In J. Adam Carter, E. Gordon, & B. Jarvis (Eds.), *Knowledge first: Approaches in epistemology and mind, chapter 11*. Oxford University Press.
Kelp, C., & Simion, M. (2017). Commodious knowledge. *Synthese, 194*, 1487–1502.
Kelp, C., & Simion, M. (2020). The C-account of assertion: A negative result (with Chris kelp). *Synthese, 197*, 125–137.
Kelp, C., & Simion, M. (2023). What is trustworthiness? (with Chris kelp). *Nous, 57*, 667–683.
Lauden, L. (2011). The rules of trial, political morality and the costs of error: or, is proof beyond a reasonable doubt doing more harm than good? In L. Green & B. Leiter (Eds.), *Oxford studies in philosophy of law*. Oxford University Press.
Lord, E., & Sylvan, K. (2021). Suspension, higher-order evidence, and defeat. In J. Brown & M. Simion (Eds.), *Reasons, justification, and defeat*. Oxford University Press.
Mace, L. (forthcoming). Doubting in philosophy. In S. Goldberg, & M. Walker (Eds.), *Philosophy with attitude*. Oxford University Press.
Millikan, R. (1984). *Language, thought, and other biological categories*. MIT Press.
Miracchi, L. (2015). Competence to know. *Philosophical Studies, 172*, 29–56. https://doi.org/10.1007/s11098-014-0325-9
Miracchi, L. (2017). When evidence isn't enough: Suspension, evidentialism, and knowledge-first virtue epistemology. *Episteme, 16*(4), 413–437.
Neander, K. (1991). The teleological notion of 'function.' *Australasian Journal of Philosophy, 69*, 454–468.
Newton, L. (2022). Epistemic anxiety and epistemic risk. *Synthese, 200*, 324.
Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy, 109*(3), 247–279. https://doi.org/10.5840/jphil201210939
Pritchard, D. (2015). Risk. *Metaphilosophy, 46*(3), 436–461.
Pritchard, D. (2016). Epistemic risk. *The Journal of Philosophy, 113*(11), 550–571.
Pritchard, D. (2018). Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence, 9*(1), 108–119.
Pritchard, D. (2022). Varieties of epistemic risk. *Acta Analytica, 37*, 9–23.
Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research, 57*, 907–919.
Simion, M. and Kelp, C. (2018b). How to be an anti-reductionist. *Synthese*.
Simion, M. (2023). Knowledge and disinformation. *Episteme*. Special Issue *Episteme* conference.
Simion, M. (2016). Perception, history and benefit. *Episteme, 13*, 61–76.
Simion, M. (2018). The explanation proffering norm of moral assertion. *Ethical Theory and Moral Practice, 21*(3), 477–488.
Simion, M. (2019a). Knowledge-first functionalism. *Philosophical Issues, 29*, 254–267.
Simion, M. (2019b). Saying and believing: The norm commonality assumption. *Philosophical Studies, 176*, 1951–1966.
Simion, M. (2019c). Hermeneutical injustice as basing failure. In P. Bondy & J. Adam Carter (Eds.), *Well founded belief: New essays on the epistemic basing relation*. Routledge.
Simion, M. (2021a). Knowledge and reasoning. *Synthese, 199*, 10371–10388.
Simion, M. (2021b). Blame as performance. *Synthese, 199*, 7595–7614.
Simion, M. (2021c). Scepticism about epistemic dilemmas. In K. McCain, S. Stapleford, & M. Steup (Eds.), *Epistemic dilemmas: New arguments, new angles*. Routledge.
Simion, M. (2021d). *Shifty speech and independent thought. Epistemic normativity in context*. Oxford University Press.

Simion, M. (2024). *Resistance to evidence*. Cambridge University Press.
Simion, M., & Kelp, C. (2018a). The constitutive norm view of assertion (with Chris kelp). In S. Goldberg (Ed.), *Oxford handbook of assertion.* Oxford University Press.
Simion, M., & Kelp, C. (2023). Trustworthy artificial intelligence. *Asian Journal of Philosophy, 31*, 447–464.
Sosa, E. (2021). *Epistemic explanations*. Oxford University Press.
Thomson, J. J. (1986). Liability and individualised evidence. *Law and Contemporary Problems, 49*, 199–219.
Williamson, T. (2005). Knowledge and scepticism. In F. Jackson & M. Smith (Eds.), *The Oxford handbook of contemporary philosophy.* Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.