http://eprints.gla.ac.uk/211491/

Deposited on 04 March 2020

De Minimis Normativism: A New Theory of Full Aptness

Full aptness is the most important concept in performance-based virtue epistemology. The structure of full aptness, in epistemology and elsewhere, is bi-levelled. At the first level, we evaluate beliefs, like performances, on the basis of whether they are successful, competent, and *apt* – viz., successful *because* competent. But the fact that aptness *itself* can be fragile – as it is when an apt performance could easily have been *inapt* – points to a higher zone of quality beyond mere aptness. To break in to this zone, one must not merely perform aptly but also in doing so safeguard in skilled ways against certain *risks to inaptness*. But how must this be done, exactly? This paper has two central aims. First, I challenge the credentials of mainstream thinking about full-aptness by raising some new and serious problems for the view. I then propose a novel account of full aptness – what I call *de minimis normativism* – which keeps all the advantages of the canonical view, avoids its problems entirely, and offers some additional payoffs.

1. Golden State Warriors point guard Stephen Curry has been working on his free throws all evening in Oakland Arena. He's hoping to finish the year with a 92% free-throw average. It's 11:30pm, and all the fans have long gone home, but he keeps shooting. He's made 47 in a row and is feeling good, though only a few janitors are present to witness it.

Around midnight, Curry hits his stride. He's at a personal-best 112 consecutive made free throws. He squares up confidently for number 113, which would be his longest streak.

Unbeknownst to Curry, an escaped beaver has been chewing for hours in pitch darkness on a utility pole just outside Oakland Arena. The beaver has nearly worked its way through the entire pole. Just one more bite, and the utility pole will drop. If the beaver's next bite is on the right side of

the pole, it will fall to the left, and it will remain propped up by a concrete barrier. However, if the beaver's next bite is on the *left* side of the pole, it will come crashing down to the right, along with the attached electrical power lines, which will cause an entire blackout inside Oakland Arena. And unfortunately for Curry, he misses more than he makes when shooting in pitch darkness.

As Curry takes a few breaths before shooting, the beaver eyes the left side of the utility pole and opens its mouth, exposing sharp orange teeth. But the beaver suddenly changes its mind. It sinks its teeth instead into the *right* side of the pole, which then falls left into the concrete barrier. Steph – aided by excellent lights in Oakland Arena – releases the shot and makes his 113th in a row.

2. A key idea in recent virtue epistemology – one that has been championed by Ernest Sosa (e.g., 2007, 2015, 2017) – is that beliefs are a species of *performance*.[1] On this way of thinking, we should normatively assess beliefs in the same way we assess performances more generally. Any performance with an aim can be evaluated along three distinct dimensions: *success* (e.g., whether it hits it aim), *competence* (whether it is produced by the exercise of a competence to attain its aim) and *aptness* (e.g., whether it is successful *because* competent). On this three-fold picture of performance assessment, *knowledge* is type-identical to apt belief, a belief that is successful (i.e., true) *because* competent.

A performance, like a belief, can be apt without being *fully apt*. Full aptness is the 'gold standard' – the fully desirable status for performances in general.[2] According to Sosa:

> *Fully apt performance*: A performance is fully apt if and only if it is guided to aptness through the agent's reflectively apt risk assessment (Sosa 2015, 69).

[1]See Fernandez (2016) for a recent collection of papers engaging with performance-theoretic virtue epistemology. Other proponents – besides Sosa – of performance-based virtue epistemology include Turri (2011, 2016), Kelp (2013), [Redacted], and with caveats, Greco (2010). C.f., Chrisman (2012) for criticism of the project of modelling belief as a kind of performance with an aim.

[2]For a more recent development of this idea, see Sosa (2017). Cf., C. Kelp et al. (2017) for a criticism of Sosa's account of full aptness.

Getting a grip on what fully apt performance (and by extension, fully apt belief – the *epistemic* gold standard) involves requires first getting a handle on what *reflectively apt risk assessment* involves. What exactly needs taken into account, and how?

Let's return now to Steph Curry, the night-time gymnasium, the utility pole and the beaver. Does Curry's shot – fired around midnight while the beaver is choosing its next (and final) bite – achieve the status of full aptness?

It would obviously *not* achieve such a status if the kind of risk assessment that is relevant to performing fully aptly must involve taking into account *all* risks to a performance's aptness. The beaver poses such a risk, one that materialises in close near-by worlds. Does the fact that Curry was *oblivious* to the beaver and its potential effects on the ambient lighting taint the quality of his shot when it flies through the hoop in perfect lighting?

More generally: What kinds of specific things must a shooter be able to competently take into account, and which things can the shooter *non-negligently ignore* without subjecting herself to credit-reducing luck when such things happen, beyond her ken, to go right?

3.   Sosa has an answer to all of this (spoiler: his view implies that Steph's shot *is* fully apt). And it's one that he thinks carries over neatly to the epistemic domain. Crucial to Sosa's answer is his new theory of *background conditions*, a theory that has emerged only in his most recent virtue episte-mology.[3]

Unfortunately, this appeal to background conditions generates two central problems for any viable account of full aptness, what I call the *arbitrariness problem* and the *demandingness problem*. Taken together, these problems expose a dilemma: attempting to accommodate the arbitrariness problem from within the canonical theory simply makes it more difficult to then deal satisfactorily with the demandingness problem.

The way forward, I'll argue, is a different and more principled approach to the kind of problem Sosa appeals to background conditions in order to solve. I develop such a proposal, which draws from tools in other areas of philosophy – and specially, from the theory of social norms and the litera-ture on *de minimis* risks. The payoff is a better account of full aptness, one that not only avoids the dilemma facing Sosa but which also clarifies more

---

[3]See in particular Sosa (2017 Ch. 13).

sharply how it is that – to perform, as well as to believe fully aptly – we must take into account certain risks while we may non-negligently ignore others.

4. A performance is *apt* if and only if it is successful *because* competent. Competences are, themselves, dispositions of an agent to perform well in a given domain of endeavour. Crucial ly, on Sosa's view, competences have a 'triple-S' constitution – seat, shape, and situation – with reference to which three kinds of dispositions can be distinguished: the innermost competence (seat), the inner competence (seat + shape), and the complete competence (seat + shape + situation).

To make this idea more concrete, consider the following illustrative example Sosa offers concerning driving competence:

> With regard to one's competence in driving, for example, we can distinguish between (a) the innermost driving competence that is seated in one's brain, nervous system, and body, which one retains even while asleep or drunk; (b) a fuller inner competence, which requires also that one be in proper shape, that is, awake, sober, alert, and so on; and (c) complete competence or ability to drive well and safely (on a given road or in a certain area), which requires also that one be well situated, with appropriate road conditions pertaining to the surface, the lighting, etc. The complete competence is thus an SSS (or an SeShSi) competence (Sosa 2017, 191–2).

Sosa alternatively uses the term 'skill' to refer to what you have when you possess the innermost competence to do something, viz., the 'seat' of a competence. And this is something we test for by asking whether you would perform reliably enough *when* in appropriate shape and while appropriately situated.[4] It accordingly doesn't count in any way against your skill to drive a car well if you would perform poorly if in improper shape (e.g., drugged, asleep, etc.) or placed in an abnormal situation – viz., extreme icy conditions or unusually high winds.

With this in mind, we can now clarify the idea that apt performances are apt when successful because competent. When a performance is apt, the

[4]For a helpful discussion of this idea in terms of what Sosa calls 'trigger manifestation conditionals', see Sosa (2010).

relevant success must issue (non-deviantly) from not merely the exercise of a skill in any kind of condition, but from a *complete* competence – viz., from a skill exercised while the subject is in proper shape and properly situated.

The matter of what *counts* as 'proper shape' and 'properly situated' matters not only for assessments of competence and thus aptness. It also has ramifications downstream for Sosa's theory of *full* aptness. Sosa's answer, in short, is that (i) the standards for 'proper shape' and 'properly situated' vary from domain to domain; and (ii) are determined by what we *care* about across particular domains of performance – viz., what the circumstances are under which good performance is valued in those domains.[5] In certain formalized contexts, these conditions are made explicit. For example, in golf, being properly situated requires possession of a putter; it also *excludes* flood conditions and storms. In the domain of salmon fishing, neither of these is the case. Relatedly, just as what counts as the shape and situation relevant to good performance varies across performance domains, so does the level of *reliability* one must manifest (when in proper shape and properly situated) to possess the relevant skill. One lacks skill at shooting free-throw if one is successful just under 50% of the time when in proper shape and properly situated; but this success rate is *above* the reliability threshold required to count as possessing skill at hitting baseballs when in proper shape and properly situated.

5. Let's now zero in more closely on *full* aptness, by running a variation on our Curry case – though this time let's leave the beaver out of the story and replace it with some lactic acid buildup. Suppose Steph continues shooting well into the early morning hours. His lactic acid is building up, but he pushes through, fighting muscle fatigue.

Stipulate that there is some threshold of lactic acid buildup after which Steph's muscles will fail to support a reliable enough free-throw-shooting form. If the threshold is crossed, it's still possible he'll make the shot. But in doing so, his shot will not be *apt.*[6] Now suppose further that Steph's shot

[5]See Sosa (2010, 466–7) and (2017, 195).

[6]This is a twist on a different case Sosa offers (e.g., Sosa 2015, 72) to illustrate first-order safety and second-order unsafety. In the case Sosa opts for, we are to imagine that the shooter is just within his threshold for reliability, but that that had the shooter been standing just an inch or so outside of where they are in fact standing, they'd fail to be reliable enough.

– taken in the wee hours in the morning – remains *within* this threshold, but just barely.

Such a shot will be apt, even though there is an important sense in which the aptness is 'fragile'. The successful shot issues non-deviantly from a complete competence. But very easily it could have not done so. Steph is in proper shape; but very easily he could not have been.[7]

What Steph's shot manifests is competence but not *reflective competence.* A shot is reflectively competent if and only if it corresponds to a competent second-order awareness *that* the shot would be apt (Sosa 2015, 68).[8] Steph lacks any such awareness in this case. While he is manifesting his first-order complete competence to shoot, he is *not* in doing so manifesting any higher-order capacities to monitor for risk of inaptness. He, after all, would easily have shot even had his lactic acid crossed the threshold for compromising his form and reliability.

Important here is that higher-order assessments of *risk of inaptness* (e.g., assessments of whether conditions are suitably favourable, whether you are in proper shape, etc.) are performances in their own right. One could be more or less reliable at doing *this* and *regardless* of whether one has reliable shooting form.

Put another way: you can have (i) aptness without reflective competence and (ii) reflective competence without aptness. Steph shooting successfully in the wee hours of the morning *just within* his threshold of reliability (while oblivious *that* he is just within this threshold) is an example of the former.[9] In the latter case, one might competently judge that a shot would be apt and thus safe enough to be worth taking. Yet it might be one

---

[7]See Sosa (2007 Ch. 2) for related cases involving aptness without safety.

[8]It's contentious whether 'awareness' is too strong here, by Sosa's own lights. As C. Kelp et al. (2017) have pointed out, by this kind of formulation, it would seem to follow that full aptness implies first-order safety. However, this is at tension with Sosa's explicit allowance of certain risks to inaptness can be non-negligently ignored by a fully apt performer. Sosa elsewhere however uses the weaker locution 'assessment', which arguably avoids this tension.

[9]Compare with Sosa's (e.g., 2015, 68) case of a huntress who shoots at a rabbit in twilight fog. The huntress believes when shooting that there is little chance of success (in light of her having drunk much wine, and the twilight fog of the conditions), and in thinking this, 'underestimates her prowess'. But she takes the shot anyway and hits the rabbit. And, unbenkownst to her, her shot was apt. She was reliable enough a shot even in those conditions. For discussion of this case and how it is subtly different from other cases Sosa uses to make a similar point, see [Redacted].

of those times when the shot misses anyhow, so that the shot is reflectively competent without being apt.[10]

Full aptness requires not only that a shot be apt *and* reflectively competent, but also that the shot be *guided to aptness* through the agent's reflectively apt risk assessment. To appreciate the relevance of the 'guidance' proviso, just suppose that Steph (in our lactic acid variation) – in consultation with doctors and Nate Silver at http://fivethirtyeight.com – gains deep and fine-grained insights into his own reliability as a shooter and how this reliability is threatened in various ways. During these consultations with doctors and Silver, he learns how to reliably spot tell-tale signs of lactic acid buildup *just at the point* that it would affect his muscles so as to compromise the reliability of his form. Now, suppose that, as he's about to shoot, Steph – through a very well-calibrated second-order risk assessment – appreciates *that* he is within, and not outside, his threshold for sufficient reliability, and thus that the shot, if taken, would be apt. Rather than to shoot, though, he flips a coin to decide *whether* to take the shot. The coin lands 'heads', and so he shoots, aptly making the shot.

In the above case, the shot is both apt *and* reflectively competent, but it is not *fully apt*; this is because the shot is not guided to aptness by the reflectively competent risk assessment conducted, but rather, by a coin flip.

6. The reader at this point might sense a tension. That's because there *is* one. I've indicated already that Sosa takes Steph's shot in the original beaver version of the case to be *fully apt*. It is fully apt even though he's oblivious to the fact that the aptness of his shot could have easily been spoiled through an electric failure that luckily, beyond his ken, didn't occur. And yet, Sosa's account of full aptness is *also* supposed to generate the result that Steph in the lactic acid variation of the case does *not* perform fully aptly, when what he's oblivious to is that the aptness of his shot could easily have been spoiled through his being in improper shape which luckily, beyond his ken, didn't transpire. That Sosa's view treats these cases disanalogously calls for explanation:

---

[10]See Sosa (2015 Ch. 3).

	shot apt?	is aptness fragile?	is performer oblivious to fragility of aptness?	shot fully apt?
Curry (beaver version)	yes	yes	yes	yes
Curry (lactic acid version)	yes	yes	yes	no

So *why* the difference of treatment, then? Here's where we get to the heart of things, and to the theory of background conditions. According to Sosa, there is an important distinction *within the class of things that could cause a performance to fail*, between:

  (i) the kinds of things a fully apt performer must heed in order to safeguard against credit-reducing luck; and
 (ii) the kinds of things he or she is free to *non-negligently* assume are already in place.

Let's look at the first category. As Sosa puts it, an athlete, in order to meet the reflective competence condition on fully apt performance:

> […] needs to consider various shape and situation factors: how tired he is, for example, how far from the target, and so on, for the many shape and situation factors that can affect performance (Sosa 2017, 191).

The lactic acid version of the case is, thus, a 'category 1' case on this characterisation. Even if Steph shoots aptly, his shot isn't fully apt if he easily could have shot inaptly *because* he easily could have been too tired. Tiredness, distance from the target, etc., are factors pertinent to *basketball*.

Let's now look squarely at Category (ii):

> But there are many factors that he need not heed. It is no concern of an athlete as such whether an earthquake might hit, or a flash tornado, or a hydrogen bomb set off by a maniac leader of a rogue state, and so on. As an athlete, he is not negligent for ignoring such factors (2017, 191).

And such things are of 'no concern' to the athlete, as such, even though earthquakes, tornadoes, bombs – as well as power outages – are the sort of things that *could* spoil a performance if they in fact materialised.

Category (ii) – viz., the kinds of things an apt performer can *non-negligently* assume are already in place – corresponds to what Sosa is calling *background conditions.* There is no simple definition of background conditions. As I've argued in previous work[11], we can identify in Sosa's recent epistemology four key characteristics of background conditions: logical, functional, epistemological, and normative.

- *Logical*: Background conditions are entailed by the presence of pertinent seat, shape and situation conditions; they must hold if the relevant 'S' [seat/shape/ situation] is in place at the time of the performance.
- *Functional*: Background conditions are orthogonal to a performance qua the kind of performance it is.
- *Epistemological*: When a background condition holds, the fully apt performer need not know that it does.[12]
- *Normative*: The quality of a performance is not reduced or in any way effected by reducing the safety a background condition.

The logical condition is the most useful for identifying which conditions are background conditions. The other characteristics make explicit properties ascribed distinctively to background conditions.

With respect to the logical condition, just consider that the shape pertinent to a figure skating competence, e.g., to be alert, awake, etc., entails whatever is necessary to support it – viz., among other things, a properly functioning thalamus. The situation pertinent this kind of competence includes a sufficiently flat ice rink; necessary for the existence of this rink is that a large sinkhole under the ice rink does not materialise. (Even more dramatically, the existence of the *earth* is necessary to support the obtaining of *any* seat/shape/situation conditions of interest to human performances of any kind).

And this leads right to the functional point: facts about the obtaining *of conditions entailed by* the presence of seat, shape, and situation conditions

---

[11]See [Redacted].
[12]Sosa (2017, 218).

– e.g., including facts about whether these conditions hold safely – seem to be completely *orthogonal* to the quality of, say, an ice skating performance, *as such.* They have nothing to do with what counts as good skating. And – here are the epistemological and normative points – because they *don't*, the *quality* of (say) a skating routine to Carmen that includes three triple axels isn't spoiled in any way if the skater who lands these jumps flawlessly was oblivious to the fact that they nearly had a sudden aneurysm while skating which would have wiped out the 'thalamus possession' background condition to being in proper shape, or that they were oblivious to the near possibility of a sink-hole under the rink. Or, for that matter, the near possibility of a beaver compromising the nearby power grid.

In sum: there is a principled reason for treating the beaver and lactic acid versions of the Curry case differently, *vis-à-vis* full aptness, despite the similarities. Though both scenarios described feature fragile aptness, what renders the aptness fragile is something that a fully apt performer can non-negligently ignore (viz., the unsafety of a background condition) in the beaver case, while what renders the aptness fragile in the lactic acid case is *not* the unsafety of a background condition, and so not something a fully apt performer can non-negligently ignore.

7.   I want to now argue that the foregoing diagnosis is problematic, even if the *verdicts* are not.

What I'm calling the *arbitrariness objection* exploits a problem with Sosa's specification of background conditions, which is that the specification doesn't screen off possible *overlap* between (i) what's entailed by the presence of the pertinent seat, shape and situation conditions; and (ii) the seat, shape and situational conditions themselves, the safety of which a fully apt performer must competently monitor.

To make this worry more concrete, consider that the presence of normal atmospheric pressure gets ruled in as a background condition on Sosa's view for almost any athletic performance type, as one cannot be in proper shape without ambient atmospheric pressure, itself a necessary condition for the presence of breathable oxygen. As such, it ought to be that a basketball player could shoot fully aptly while ignoring threats to atmospheric pressure – even when such threats are modally close.

But *dips* in atmospheric pressure are well known to – like the build up of lactic acid! – also lead to one's shape being compromised, e.g., by bringing

about joint stiffness. When thinking about atmospheric pressure in this way, it looks as though a fully apt performer could non-negligently ignore nearby threats to normal levels of atmospheric pressure *only if* they can also non-negligently ignore more mundane threats to being in proper shape, including lactic acid buildup, pains in one's muscles, etc. But these are exactly the kinds of things Sosa says a fully-apt performer *can't* be oblivious to.

The atmospheric pressure case raises a question: in cases of 'overlap', which factor trumps the other? Sosa's view has no straightforward answer here. And this problem itself compounds when we consider that overlap cases will be relatively common. To take another example, consider the way professional chess players monitor their glucose levels through nutrition during games. The kind of shape appropriate to elite chess includes alertness and mental acuity, the obtaining of which *entails* normal glucose levels of the very sort it is (within the standards of professional chess) taken to be negligent *not* to monitor – especially in classical format games which may last 6 hours.[13]

A consequence of these kinds of overlap cases is a kind of arbitrariness. Wherever there is overlap, in the absence of an adjudicating rule, assessments that overlooking certain risks to aptness is non-negligent will be arbitrary assessments. This arbitrariness exposes a problem with Sosa's theory of background conditions, and it is a problem that the theory of full aptness unavoidably inherits.

8. Any solution to the above problem will require, at some point or another, taking a principled stance on what kind of risks to the aptness of a performance, for a given performance type, *can't* be non-negligently ignored when one performs fully aptly. On this score, the best clue we find from within Sosa's virtue epistemology is his remark that 'it is no concern of *an athlete as such* whether an earthquake might hit […].'[14] Perhaps a clearer theory of what, for any given performance type, is of concern to that performer of that performance type *as such* would offer a way forward. But if such a strategy is going to work specifically for a theory of fully apt

[13]See, for instance, recent studies reported by Alifirov, Mikhaylova, and Makhov (2017).

[14](Sosa 2017, 191).

judgment – as a theory of the gold standard in epistemology – we'd need to know what is of concern to an *inquirer as such*?

But what is of concern to an inquirer *as such*? On one very reasonable way of answering this, the answer is: in principle, *everything*![15] And this marks important sense in which inquirers *as such* differ from, say, basketball players as such, the latter of whose concerns exclude all but very specific kinds of things. And this all looks like bad news for the prospects of dealing with the arbitrariness objection from within Sosa's epistemology in a way that allows the theory of full aptness to have applicability in epistemology. All indications are that the consequent theory of full aptness would tolerate almost *no* negligence that didn't count as tainting the quality of the relevant intellectual performance. Call this the *demandingnes problem.*

The arbitrariness problem and the demandingness problem pull Sosa's view in opposite directions. The arbitrariness problem – generated by the theory of background conditions – looks resolvable only with reference to a clear account of what kinds of risks to the aptness of a performance, for a given performance type, *can't* be non-negligently ignored. I've shown that, working from within Sosa's theory at least, the most promising kind of account is one that is transferable to epistemology only at the cost of incurring the demandingness problem.

9. In what follows, I want to propose a different way forward – one that sidesteps the above dilemma by incorporating into a theory of full aptness some insights from the (i) the theory of social norms; (ii) and the literature on *de minimis* risks.

Here is the core statement of view I will now defend, what I call *de minimis normativism*:

> (†) A fully-apt performer can't non-negligently ignore *practice-relative risks* to the inaptness of a given performance that occurs within that practice-type, except when these risks count as *de minimis* with reference to practice-sustaining rules.

[15]A possible exception here concerns what are called 'trivial truths'. For example, long disjunctions with just one true disjunct. Iterating such disjunctions by adding further false disjuncts generates new truths, though the intellectual value of such truths as an object of inquiry is contested. For discussion, see, e.g., Kvanvig (2008).

Several pieces of terminology need clarified. The first concerns *practice-sustaining rules*. Let's define generally – in a way that abstracts from athletic and epistemic domains – a 'practice' as a way of doing things and a 'rule' as a prescriptive principle[16] or standard of conduct.[17] Prescriptive rules (hereafter, 'rules') can be primary or derivative (a distinction that we will return to).[18] To a first approximation, primary rules say 'do ϕ' or 'don't do ϕ'. For example: don't break promises.[19] Derivative rules are generated by primary rules and take the form: 'do what a person disposed to satisfy the primary rule would do'. (Example: try to bring it about that you don't break promises.) Rules are important to practices: they 'hold practices together'. But *how* do they do this?

A straightforward recent answer has been defended by John Turri (2017), one that is value driven:

> *Practice-sustaining rules*: A rule normatively sustains a practice if and only if the value achieved by following the rule explains why agents continue following that rule.

"Don't break promises" is a sustaining rule for many kinds of practices: the value achieved by following this rule explains why clergy as well as bankers continue to follow it. Yelling "bingo" if and only if you have a bingo is a practice sustaining rule just for bingo: the value of doing this explains why players of bingo keep doing this.

A practice might have many rules, though only some of these play the role of sustaining it, by leading to 'reproduction via value produced' – alternatively, by having *reproduction value*. A simple (albeit imperfect) heuristic for assessing whether a rule has reproduction value is to check whether it has derivative rules that themselves have reproduction value. If not, then this counts against the rule itself being a rule that sustains the practice, as opposed to one that merely features in the practice.

Many practices *include* performances. They do so when performances are prescribed, in certain conditions, by rules that sustain the practice. For

---

[16]For discussion on the difference between prescriptive and evaluative norms, see Simion, Kelp, and Ghijsen (2016).

[17]Here I am following John Turri (2017 Sect. 1).

[18]See Williamson (2018) and Simion, Kelp, and Ghijsen (2016) for discussion.

[19]This is the example typically used by Williamson (2018) to characterise primary norms.

example, the practice of archery includes the performance of shooting an arrow at a target. The practice of playing chess includes the performance of castling to defend the king. The practice of inquiry includes belief.

Practice-relative risks to the inaptness of a performance, within a given practice, can now be defined in terms of practice-sustaining rules as follows:

> *Practice-relative risks*: Risks to the inaptness of a performance, within a given practice ϕ, are ϕ-relative risks if and only if, were the performance inapt, it would constitute a violation of (at least one) primary ϕ-sustaining rule or rules.

Consider again basketball and the risk a shooter runs when she is *barely inside* her reliability threshold and, oblivious to this, shoots the shot aptly. In the nearest worlds where *that* shot is inapt, a primary ϕ-sustaining rule is violated. After all, these are worlds where the shooter steps just a few inches back before shooting. And it is a primary practice-sustaining rule in basketball that, *ceteris paribus,* you should not shoot too far from the basket, beyond your sufficient threshold for reliability.[20] (Note: this rule has its own derivative rules that have reproduction value – viz., 'check how far out you are before you shoot!'.) But what makes a risk to the aptness of a performance *practice-relative*, with respect to a practice ϕ, is whether, were it inapt, a primary ϕ-sustaining rule would be violated.

Of course, when there is a hungry beaver chewing on the utility pole outside the arena, there is (like when one is nearly outside her reliability threshold) *also* a risk to the inaptness of the shot. But, crucially, the risk to the aptness of the shot posed by the beaver is not a *practice-relative risk.* In the nearest world in which sudden darkness spoils the aptness of Curry's shot, it's *not* the case that Curry violates any plausible primary ϕ-sustaining rule. 'Don't shoot in the dark' is not a good candidate for such a rule. It lacks any obvious derivative rule with reproduction value. (After all, it's implausible that the value achieved by following some rule taking

---

[20]Within the basic rules of basketball (e.g., a shot counts as 2-points, you can not walk while carrying the ball) it is a practice sustaining rule that you not shoot willy nilly any time you have the ball. The value of shooting (all else equal) only high percentage shots explains why this continue to be followed. For related discussion from Sosa on the impropriety of 'blind shooting', in the context of discussing why shooters *aim* to shoot not only successfully but aptly, see Sosa (2015, 71–72).

into account *immediate darkness* possibilities explains why agents continue to follow any such rule. On the contrary: the value achieved by *ignoring* such possibilities would explain why players carry on ignoring them.)

10. For reasons that will be apparent shortly, a final piece of the puzzle is needed. This concerns the *de minimis* risk proviso that features in (†). The phrase *de minimis* derives from the Latin sentence *de minimis non curat lex*, which translates (roughly) to 'The law should not concern itself with trifles' (e.g., the crime of stealing a penny).[21] In decision theory, risks are termed *de minimis risks* whenever they are judged to be so 'small' that they should be ignored.[22] The concept is an especially important in health and environmental decision making.[23]

So why, exactly, is a proviso of this sort needed in (†)? After all, if the account *already* gets the result that full aptness isn't undermined by the nearness of the beaver/darkness possibility – given that this risk to the inaptness of the shot isn't a *practice-relative* risk – then isn't the inclusion of a *de minimis* proviso redundant?

The answer here is 'no'. Whereas there is no (in Turri's terms) 'reproduction value' achieved in basketball by attending to immediate darkness possibilities, there *is* reproduction value achieved by not shooting when too tired. Usually this is done in basketball by following derivative rules such as 'keep an eye on tell-tale signs that lactic acid is building up[24]' etc. But there are rarer sources of tiredness in basketball, viz., ingesting halothane gas emitting from a small hole in the court. Suppose that Curry unfortunately does exactly this, which causes him to become tired *without* lactic acid buildup, *almost* tired enough that his shooting form is compromised, but (beyond his ken) it's not. *Within* his sufficient reliability threshold, Curry then shoots and makes, oblivious that he was nearly too tired to shoot with reliable enough form *because* he is oblivious to the fact that he's ingested what is nearly a reliability-compromising dose of holothane gas.

[21]See Peterson (2002, 47).

[22]See Peterson (2002) for discussion.

[23]See, for example, Sandin (2005); Rulis (1986); Rhodes et al. (2011); Whipple (2012) and Mumpower (1986).

[24]Note that this is an example of a rule that one who is disposed to comply with the primary rule would aim to comply with. It is derivative because it prescribes a way of attempting to comply with the primary norm rather than prescribing simply that the primary norm be complied with.

Without some kind of *de minimis* proviso in (†), it looks like the account in (†) would problematically lump the halothane gas version of the case with the lactic-acid *rather* than with the beaver version of the case, as one where full-aptness is undermined. But this is a bad result: it seems after all that the holothane and beaver risks – the monitoring for each of which seems equally irrelevant as the other is to quality basketball – should stand and fall together, *even though* monitoring for signs of tiredness is itself prescribed by derivative practice-sustaining rules and monitoring for signs of a sudden power-outages is not.

With this in mind, the formulation of the *de minimis* proviso I want to now defend, as a component of the wider principle (†), is the following:

> *De minimis proviso* (†): For any practice ϕ, a ϕ-relative risk, X, to the inaptness of a given performance, P, is *de minimis* with reference to ϕ, if and only if the safety of P against X can't be easily increased through adherence to one or more derivative ϕ-sustaining rules.

The safety of a performance against a risk concerns how easily the risk would materialise.[25] Doing something to *increase* the safety of a performance against a risk is to do something that makes it less easy (holding fixed that you've done that thing) that the risk will materialise – viz., that, holding fixed that you've done that thing, the risk event materialises in further-out worlds than before.[26] The *de minimis* proviso above says that a risk is *de minimis* when, specifically, adhering to derivative rules of the relevant practice is *not* among the things that can easily increase the safety of a performance against a practice-relative risk.

Two qualifications are in order here. First, regarding the 'derivative' qualifier. Here is why it matters. Remember: a risk is practice-relative only if, were the performance inapt, it would constitute a violation of (at least one) primary ϕ-sustaining rule or rules. In the holothane gas case, the primary practice-relative rule that would be violated in the nearest worlds where the shot is inapt is "don't shoot when too tired", a rule with clear reproduction value in basketball. Now, *trivially* one can increase the performance's safety against that risk by adhering to *that* primary rule – viz., by

---

[25]For some representative discussions of safety in epistemology, see Rabinowitz (2011), Pritchard (2005), Comesaña (2005), and Ballantyne (2012).

[26]For related discussion, see Pritchard (2016).

not shooting when too tired. What one can't do, however, is easily increase the safety of the performance against that risk by adhering to *derivative* ϕ-sustaining rules. This is because adhering to derivative rules that have reproduction value in basketball (e.g., check for familiar signs that one is too tired) doesn't *easily* increase the safety of the performance against the risk to inaptness that is posed by holothane gas; monitoring in *those* ways, you'd never see it coming – at least, by following such rules. You could, by contrast, safeguard swimmingly against the holothane gas risk by toting around a gas monitor and checking it regularly while on the court. But attempting to comply to the primary rule 'don't shoot when too tired' by adhering to *this* derivative rule blatantly lacks reproduction value in basketball.

Second, regarding the '*can't be easily* increased' locution. Why not just: '*can't be* increased?' Consider that one who gets very lucky could increase the safety of the shot against the holothane gas risk to its inaptness by adhering to a derivative rule with reproduction value, like 'check for familiar signs that one is too tired'. As it happens, one of the *less common* side effects of halothane gas is difficulty breathing, a side effect that overlaps with one of the tell-tale signs of lactic acid buildup as when one typically becomes tired. One *could*, as it were, 'get lucky' and experience this rarer symptom and correctly identify it as a marker of tiredness. In doing so, the safety of the performance against the holothane gas risk to inaptness *would* be increased through adherence to one or more derivative ϕ-sustaining rules. It just wouldn't be *easily* increased.

With the *de minimis* proviso now fully unpacked, it should be clear how it neatly separates the lactic acid and holothane gas cases, by ruling the latter in as *de minimis*, and the former out. And this means that the wider account – *de minimis* normativism – rightly classifies the holothane gas and beaver cases as both cases where a performance is (unlike in the lactic acid case) fully apt, and despite the fact that holothane gas poses a practice-relative risk to the aptness of the shot, and the beaver does not.

11. Let's take stock. There is an important distinction *within the class of things that could cause a performance to fail*, between: (i) the kinds of things a fully apt performer must heed in order to safeguard against credit-reducing luck; and (ii) the kinds of things he or she is free to *non-negligently* assume are already in place.

What makes the difference between (i) and (ii)?

- Sosa's answer: A fully apt performer can't non-negligently ignore risks to the inaptness of a performance unless those risks are due to the unsafety of background conditions.
- My answer (viz., *de minimis* normativism): A fully-apt performer can't non-negligently ignore practice-relative risks to the inaptness of a performance that occurs within that practice-type, except when these risks count as *de minimis* with reference to practice-sustaining rules.

I've argued that Sosa's answer generates a dilemma. In particular, it was shown that his view is susceptible to a kind of arbitrariness problem, and further, that this arbitrariness problem can't be solved on Sosa's own terms without exacerbating an independent problem – one to do with demandingness – when the proposal is transferred over to the epistemic case, viz., as an account of fully apt judgment, specifically.

Moreover, we saw *why* it is that the answer Sosa gives leaves him open to the arbitrariness problem in the first place. This is because, by giving the answer he does above, he fails to screen off possible *overlap* between (i) what's entailed by the presence of the pertinent seat, shape and situation conditions; and (ii) the seat, shape and situational conditions themselves, the safety of which a fully apt performer must competently monitor. Two cases were raised to make this point about overlap: one to do with atmospheric pressure and another to do with chess.

My proposal can deal with both of these cases straightforwardly, and with out any reference to background conditions. On my view, an apt chess move that could easily have been inapt because the player (oblivious to her crashing glucose levels) could easily have been in improper shape is *not* fully apt. This is because on my view a fully apt performer can't non-negligently ignore practice-relative risks to the inaptness of a performance, and *this* is an instance of practice-relative risk. After all, were the performance inapt, there would be a violation of a primary practice-sustaining rule or rules – viz., "don't play with compromised mental acuity" a primary rule that generates derivative rules with reproduction value, such as "try to keep acuity sharp by monitoring glucose levels." Moreover, on my view, this practice-relative risk is not *de minimis* because monitoring for glucose levels in this case *would* increase the safety of the performance against inaptness.

It is important to note that the above rationale does *not* imply that, to make a chess move fully aptly, the chess player must be thinking throughout the game about about glucose levels. Doing *that* would not have reproduction value in chess, even if it would increase the safety of the performance against inaptness from mental sluggishness. (Compare: monitoring not only for regular signs of tiredness but for signs one has ingested halothane gas would not have reproduction value in basketball even if it would increase the safety of a performance against inaptness from tiredness when there happens to be holothane gas about.)

Likewise, the proposal can diagnose the atmospheric pressure case – the other case ruled as an 'overlap' case on Sosa's view – in a principled way. Though, in this case, the risk (at least in basketball – will be *de minimis*, unlike in the chess case. The thinking here is as follows: *even though* dips in atmospheric pressure can lead one's one's shape being compromised, e.g., by bringing about joint stiffness, this is a *de minimis* risk for a basketball player, one she can shoot fully aptly while ignoring. This is because there is no derivative basketball sustaining rule (viz., no derivative rule *with reproduction value in basketball*) the adherence to which would increase the safety of a basketball shot against *that* risk. One could safeguard against it, no doubt – perhaps by being mindful of both the forecast and ways in which the arena could open up to the elements outside. But there is reproduction value in basketball to simply ignoring such risks, rather than to concern oneself with them.

12.   Let's turn now to epistemology specifically. Human knowledge – viz., the kind of knowledge that adult humans aspire to – is not just apt belief[27], but *fully apt* belief.[28] Fully apt believers can, like any kind of fully apt performer, non-negligently ignore certain risks to the inaptness of beliefs they hold aptly. My preferred view – *de minimis normativism* – has clear implications for just *which* risks to inaptness these are, implications that I've shown contrast favourably with those of Sosa's.

[27]This is what Sosa terms 'animal knowledge'. For discussion, see Sosa (2007 Ch. 2).

[28]This idea has its roots in Descartes' distinction between *cognitio* and *scientia*, where the ascent from lower (*cognitio*) to higher knowledge (*scientia*) marks an intellectually valuable transition attained through reflection on one's epistemic position. Descartes distinction has close parallels between, on the performance-theoretic framework, apt and fully apt belief. For detailed discussion of this parallel, see Sosa (2017 Ch. 1) and Reed (2012). For criticism, see for example Kornblith (2004) and Fumerton (2004).

In this section, I want to highlight a further advantage of the view, one that pertains to the place of epistemic *blame* in an account of full aptness.[29] To this end, let's consider two different twists on Sosa's kaleidoscope perceiver case (Sosa 2007, 31).

> KALEIDOSCOPE-1: You see a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented you with a red-light+white-surface combination as with the actual white-light+red-surface combination. You have been cautioned – rightly – that there are jokesters about, and you look carefully for signs of a jokester or a lighting machine, but spot nothing. The jokester, which *fortunately* gave you the white-light+red-surface combination rather than the red-light+white-surface, was, *unfortunately*, well hidden. Oblivious to this, and thus to the nearness of the red-light+white-surface possibility, you believe the surface is red.

> KALEIDOSCOPE-2: You see a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented you with a red-light+white-surface combination as with the actual white-light+red-surface combination. You have been cautioned – rightly – that there are jokesters about. You give a quick, half-hearted scan for signs of a jokester or a lighting machine, but spot nothing. The jokester, which *fortunately* gave you the white-light+red-surface combination rather than the red-light+white-surface, was, *unfortunately*, well hidden. Oblivious to this, and thus to the nearness of the red-light+white-surface possibility, you believe the surface is red.

Here are three observations about these cases:

[29] For a recent general discussion of epistemic blame, see Brown (2017).

(a) *Both* KALEIDOSCOPE-1 and KALEIDOSCOPE-2 feature apt belief.[30]

(b) *Neither* KALEIDOSCOPE-1 nor KALEIDOSCOPE-2 features fully apt belief because neither features reflectively apt risk assessment.

(c) The belief formed in KALEIDOSCOPE-1 is *better* than the belief formed in KALEIDOSCOPE-2.

Sosa's account of full aptness treats the two versions of the case the same (viz., as apt belief that falls short of full aptness) and so can explain observations (a) and (b), but not (c). Mine, by contrast, has a simple explanation for why the belief in KALEIDOSCOPE-1 is better than the belief in KALEIDOSCOPE-2, even though *both* are (as Sosa's view rightly holds) cases of apt belief that fall short of full aptness. This is because, on my proposal, the apt believer in KALEIDOSCOPE-1 falls short of full aptness *blamelessly* while the apt believer in KALEIDOSCOPE-2 does not.

A common way of thinking about blameless norm violations – following Williamson (2018) – is as follows: one violates a primary prescriptive norm (e.g., keep your promises) if and only if one fails to comply with that norm.[31] To assess whether the violation of the primary norm is *blameless* we ask whether one has complied with the relevant derivative norm – viz., "do what a person disposed to comply with the primary norm would do". If they have complied with this derivative norm, then even if they've not complied with the primary norm, their violation of it is blameless. If not, then it is not.[32]

In neither KALEIDOSCOPE-1 nor KALEIDOSCOPE-2 is the subject is a position to believe fully aptly while taking for granted that the light-

[30]In neither case does the jokester actually implement the bad red-light+white-surface combination. Because you are in both cases given the good white-light+red-surface combination, you are, as Sosa puts it, 'exercising your faculty of color vision in normal conditions of lighting, distance, size of surface, etc., in conditions generally appropriate for the exercise of color vision' (2007, 31) and consequently your true belief issues from the exercise of a complete competence.

[31]For a discussion of blame and criticism with respect to norm violations in action and assertion specifically, see C. Kelp and Simion (2017).

[32]Note that Williamson himself uses the notion of blameless norm violation in the service of a specific dialectical objective, which is to argue that victims of BIV scenarios lack justification. For the present purpose I am not taking a stand on how epistemic justification maps on to epistemic norm compliance. Though see Simion, Kelp, and Ghijsen (2016) for a helpful recent discussion on this issue.

ing is normal as opposed to the manifestation of a jokester's trick. And this is because both subjects have a testimonial basis for thinking, rightly, that the situation could very easily be improper (red-light+white-surface combination) for forming beliefs about the colour of the wall. Both the subjects in KALEIDOSCOPE-1 and KALEIDOSCOPE-2 should accordingly form wall-colour beliefs only if the red-light+white-surface combination is ruled out. (It is, after all, a good rule of inquiry to rule out known-to-be nearby error possibilities![33]) The *derivative* rule in play here, one with clear reproduction value in the practice of inquiry, is to do what one who is disposed to rule out the relevant error possibilities would do. This is done in KALEIDOSCOPE-1, *even though* the jokester is never actually spotted! Similar due diligence is however *not* performed in KALEIDOSCOPE-2, where like in KALEIDOSCOPE-1, the jokester is not spotted. Though neither case features fully apt belief, the former falls short of full aptness *blamelessly* and the latter does not. This explains why the belief is better in KALEIDOSCOPE-1 than KALEIDOSCOPE-2.

13.   This essay has been principally concerned with the kind of quality that lies beyond aptness and which a performance attains only when aptness features alongside risk assessment that must itself be of a certain quality. The traditional theory of full aptness – championed by Sosa – is right to register that full aptness is compatible with monitoring for *some* but not *all* risks to inaptness. But which are those that can be ignored and why?

The most promising answer thus far, one offered by Sosa himself, appeals to a theory of *background conditions*, a theory that – despite its initial plausibility – was shown to face intractable problems that the wider theory of full aptness then inherits. In its place, I've defended a new account of full aptness – *de minimis normativism.* This view – which gives social norms a central role – sidesteps entirely the problem Sosa appeals to background conditions in order to solve. Moreover, the view has its own distinctive benefits. One such key benefit is that the view can easily explain when an apt performer falls short of full aptness blamelessly. The distinction, within the class of apt performances that fall short of full aptness, between those that are blamelessly merely apt and those which are not, is one that the traditional theory of full aptness problematically elides. Moreover, by of-

---

[33]For related discussion of relevant alternatives, in the context of fully apt judgment, see Sosa (2017, 237).
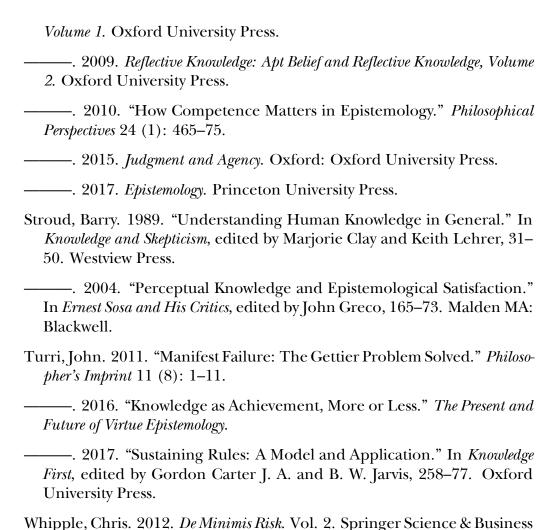
fering a clear account of what higher-order risks can be ignored and why, *de minimis normativism* offers new resources, beyond those offered by the traditional theory of full aptness, as an *anti-sceptical* strategy. The distinctive resources of bi-level virtue epistemology, generally, as an anti-sceptical epistemology is one of the theory's most well-known selling points.[34] And the advantages of *de minimis normativism* as a way of explaining how certain risks can be non-negligently ignored only strengthens the options a bi-level virtue epistemologist already has against the sceptic. But properly putting these new advantages to use against the sceptic will need to wait for another occasion.

References

Alifirov, AI, IV Mikhaylova, and AS Makhov. 2017. "Sport-Specific Diet Contribution to Mental Hygiene of Chess Player." *Theory and Practice of Physical Culture*, no. 4: 30–30.

Ballantyne, Nathan. 2012. "Luck and Interests." *Synthese* 185 (3): 319–34.

Brown, Jessica. 2017. "What Is Epistemic Blame?" *Noûs*. https://doi.org/10.1111/nous.12270.

Chrisman, Matthew. 2012. "The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions." *Journal of Philosophy* 109 (10): 588–612.

Comesaña, Juan. 2005. "Unsafe Knowledge." *Synthese* 146 (3): 395–404.

Fernandez, Miguel Angel, ed. 2016. *Performance Epistemology: Foundations and Applications.* Oxford University Press.

Fumerton, Richard. 2004. "Achieving Epistemic Ascent." In *Ernest Sosa and His Critics*, edited by Greco John, 72–85.

Greco, John. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity.* Cambridge University Press.

Kelp. 2013. "Christoph." *Australasian Journal of Philosophy* 91 (2): 265–78.

[34]See, especially, Sosa (2009). Cf., Stroud (1989, 2004).

Kelp, Christoph, Cameron Boult, Fernando Broncano-Berrocal, Paul Dimmock, Harmen Ghijsen, and Mona Simion. 2017. "Hoops and Barns: A New Dilemma for Sosa." *Synthese*, 1–16.

Kelp, Christoph, and Mona Simion. 2017. "Criticism and Blame in Action and Assertion." *The Journal of Philosophy* 114 (2): 76–93.

Kornblith, Hilary. 2004. "Sosa on Human and Animal Knowledge." In *Ernest Sosa and His Critics*, edited by Greco John, 126–34.

Kvanvig, Jonathan. 2008. "Pointless Truth." *Midwest Studies in Philosophy* 32: 199.

Mumpower, Jeryl. 1986. "An Analysis of the de Minimis Strategy for Risk Management." *Risk Analysis* 6 (4): 437–46.

Peterson, Martin. 2002. "What Is a de Minimis Risk?" *Risk Management* 4 (2): 47–55.

Pritchard, Duncan. 2005. *Epistemic Luck.* Oxford University Press.

———. 2016. "Epistemic Risk." *The Journal of Philosophy* 113 (11): 550–71.

Rabinowitz, Dani. 2011. "The Safety Condition for Knowledge." *Internet Encyclopedia of Philosophy.*

Reed, Baron. 2012. "Knowledge, Doubt, and Circularity." *Synthese* 188 (2): 273–87.

Rhodes, Rosamond, Jody Azzouni, Stefan Bernard Baumrin, Keith Benkov, Martin J Blaser, Barbara Brenner, Joseph W Dauben, et al. 2011. "De Minimis Risk: A Proposal for a New Category of Research Risk." Taylor & Francis.

Rulis, Alan M. 1986. "De Minimis and the Threshold of Regulation." *Food Protection Technology*, 29–37.

Sandin, Per. 2005. "Naturalness and de Minimis Risk." *Environmental Ethics* 27 (2): 191–200.

Simion, Mona, Christoph Kelp, and Harmen Ghijsen. 2016. "Norms of Belief." *Philosophical Issues* 26 (1): 374–92.

Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge,*

Volume 1. Oxford University Press.

———. 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume 2.* Oxford University Press.

———. 2010. "How Competence Matters in Epistemology." *Philosophical Perspectives* 24 (1): 465–75.

———. 2015. *Judgment and Agency.* Oxford: Oxford University Press.

———. 2017. *Epistemology.* Princeton University Press.

Stroud, Barry. 1989. "Understanding Human Knowledge in General." In *Knowledge and Skepticism,* edited by Marjorie Clay and Keith Lehrer, 31–50. Westview Press.

———. 2004. "Perceptual Knowledge and Epistemological Satisfaction." In *Ernest Sosa and His Critics,* edited by John Greco, 165–73. Malden MA: Blackwell.

Turri, John. 2011. "Manifest Failure: The Gettier Problem Solved." *Philosopher's Imprint* 11 (8): 1–11.

———. 2016. "Knowledge as Achievement, More or Less." *The Present and Future of Virtue Epistemology.*

———. 2017. "Sustaining Rules: A Model and Application." In *Knowledge First,* edited by Gordon Carter J. A. and B. W. Jarvis, 258–77. Oxford University Press.

Whipple, Chris. 2012. *De Minimis Risk.* Vol. 2. Springer Science & Business Media.

Williamson, Timothy. 2018. "Justifications, Excuses, and Sceptical Scenarios." In *The New Evil Demon,* edited by Julien Dutant and Fabian Dorsch. Oxford: Oxford University Press.