Justification, Normalcy and Randomness

Martin Smith

Some random processes, like a series of coin flips, can produce outcomes that seem particularly remarkable or striking. This paper explores an epistemic puzzle that arises when thinking about these outcomes and asking what, if anything, we can justifiably believe about them. The puzzle has no obvious solution, and any theory of epistemic justification will need to contend with it sooner or later. The puzzle proves especially useful for bringing out the differences between three prominent theories; the probabilist theory, the normic theory and a theory recently defended by Goodman and Salow.

I Three Stances

Suppose 100 fair coins are about to be flipped. Would I be justified in believing that they won't all land heads? There is an immediate temptation to answer 'yes'. While it's possible that we'll get a sequence of 100 heads this would surely be an incredible fluke, and not something that I need to take seriously. But now consider a nondescript sequence of 100 heads and tails – HTTHTHTTTH…. Would I be justified in believing that the coins won't land in this sequence? There is, I think, some inclination to answer 'no'. After all, we know that the coins have to land in some sequence or other and there doesn't seem to be anything that entitles me to rule this particular sequence out. But this sits uneasily alongside our answer to the first question. While a sequence of 100 heads might seem particularly remarkable it is *just as likely* to come up as the unremarkable sequence HTTHTHTTTH…. If we were to repeat this experiment over and over we would expect to observe these two sequences just as often. So why treat them differently?

There are three ways we might respond at this point. One thing we could do is switch our second answer to 'yes' and affirm that I do after all have justification for believing that the coins won't land in the sequence HTTHTHTTTH…. Since this is just an arbitrary sequence, what we are really committing to here is the claim that, for *any* possible sequence that could result from 100 coin flips, I have justification for believing that the coins won't land in that sequence. Call this the *lavish* stance – 'lavish' because, on this view, I have justification for believing quite a lot of things about the outcome of the coin flips.

Another option we have at this point is switch our first answer to 'no' – to back-pedal on the idea that I have justification for believing that we won't get a sequence of 100 heads. Since this is about as striking a sequence as we can imagine, what we are really committing to here is the claim that, for any possible sequence that could result from 100 coin flips, I *lack* justification for believing that the coins won't land in this sequence. Call this the *austere* stance. On this view I don't have justification for believing much at all about how the coins will land.

The final option is to stick with our first impressions, answering 'yes' to the first question and 'no' to the second, in spite of the fact that these two sequences are equally likely. We are then committed to the claim that, for *some* of the possible sequences that could result from 100 coin flips, I have justification for believing that they won't come up – but not for others. Call this the *differential* stance. More work will be needed to determine exactly which sequences fall into which category – but an initially attractive suggestion is that I have justification for dismissing those sequences that feature a particularly large or small number of heads, but I lack justification for dismissing the others.

This example serves as a useful testing ground for different theories of epistemic justification. In fact, we can easily find theories in the literature which will deliver each of the three aforementioned stances. The lavish stance fits with a *probabilist* theory of justification – which is perhaps the theory that epistemologists most often default to. According to this theory one has justification for believing a proposition P just in case the probability of P, given one's evidence, exceeds a high threshold. Given the evidence that the coins are fair and are going to be flipped 100 times, the 2^{100} sequences that could result are all equally likely. For any one sequence, then, the probability that it *won't* result is going to be $1 – 1/2^{100}$ which would, presumably, be higher than any threshold we would care to set. According to the probabilist theory, I have justification for believing, of each possible sequence, that it won't come up.

One theory of justification that would fit with the austere stance is the *normic* theory that I have defended in previous work (Smith, 2010, 2016, 2018a, 2018b, 2022a). According to this theory one has justification for believing a proposition P just in case, given one's evidence, the situation in which P is false would be less normal, in the sense of requiring more explanation, than the situation in which P is true. Given the evidence that the coins are fair and are going to be flipped 100 times, the 2^{100} sequences that could result are all equally normal, in that no one sequence would require any more explanation than any other. On the normic theory I don't have justification for dismissing any possible sequence. This is even true of sequences that have a remarkable pattern, like a sequence of 100 heads. We may be leery of the idea that a random process could produce such a patterned outcome, but it can of course – indeed that's part of what it is for the process to be random. As observed above, if the experiment were repeated over and over we would expect a sequence of 100 heads to come up just as often as any other sequence.[1]

The normic theory lends itself to a particular kind of formalism. Suppose the possible worlds that are consistent with one's evidence can be ordered with respect to how normal they are. If the most normal worlds in which P is false are less normal than the most normal worlds in which P is true then this encodes the fact that, given one's evidence, the falsity of P would require more explanation than its truth.[2] A justification attribution can then be treated as quantifying over a privileged set of worlds – one has justification for believing P just in case P is true in all of the most normal worlds that are consistent with one's evidence. Since the most normal worlds in which 100 fair coins are flipped include worlds in which every possible sequence comes up, I won't have justification for dismissing any possible sequence, including a sequence of 100 heads.[3]

---

[1] The 'remarkable' outcomes are the ones that might be described as *coincidences* if and when they occur. So we might say that it's an amazing coincidence for a random process to produce a run of 100 heads, but we wouldn't say that about HTTHTHTTTH…. It's difficult to lay down general criteria for when an event should count as a 'coincidence' (see, for instance, Lando, 2017) – and this is not a topic that I will go into here. But one point that may tie in with the material to come is that most analyses of coincidence involve some kind of *psychological* condition – so whether an event counts as a 'coincidence' is not a wholly objective feature of the event itself, but depends in part on our psychology (see Lando, 2017, p133).

[2] There are a few potential ways to bridge the gap between the informal conception of normalcy, in terms of explanatory priority, and the possible worlds model. The approach in Smith (2022a, section 5) is based, roughly, on the idea that propositions can be ranked according to the amount of explanation they require, from which a ranking of possible worlds can then be extracted via a standard possible worlds model of propositions. The details of this construction will not matter for present purposes.

[3] In Smith (2017) I argued that one shouldn't be *surprised* by a long run of coins coming up heads, using some considerations related to the normic theory. Consider the following principle: One should be surprised by an event just in case one had justification for believing that it wouldn't occur. If we accept this then the austere stance predicts that none of the possible outcomes of 100 coin flips should be surprising. The lavish stance, on

But if we untether this formalism from the 'explanatory' conception of normalcy that informs the normic theory, then it could also be used, in principle, to subserve the *differential* stance. If the most normal worlds in which we get a sequence of 100 heads are rated to be *less* normal than the most normal worlds in which we get the sequence HTTHTHTTTH… then the theory could easily predict that I have justification for dismissing the former but not the latter. If we are thinking of abnormality in terms of a need for special explanation then there would be no rationale for ranking the worlds like this – but the formalism itself presents no barrier to such a ranking. As we will see in the next section, this kind of model plays a role in the most systematically developed theory of justification that fits the differential stance – the theory defended by Goodman and Salow (2023).

Each of the three stances has counterintuitive consequences, and there's no easy way to tell, at a glance, which is best. Unfortunately, subjecting these stances to close scrutiny doesn't really make the decision any easier – but that, at any rate, is my aim here. More precisely, my aim is to push back against the differential stance, which has perhaps been gaining some traction in the recent literature. I will argue that some of the advantages that have been claimed for the differential stance are illusory or, at the very least, intertwined with significant costs. My tentative conclusion is that the austere stance offers the best overall combination of costs and benefits.

2 In Favour of the Differential Stance

While we can distinguish 2^{100} different outcomes that could result from 100 fair coin flips, we undoubtedly find it natural to group certain outcomes together – by, for instance, counting up the number of heads that are showing. This will result in groups of very different sizes – there are, for instance, around 10^{29} possible sequences that feature exactly 50 heads, and only one sequence that features 100. If we plot the proportion of sequences involving n heads then this will approximate a normal distribution with a mean of 50 and standard deviation of 5.

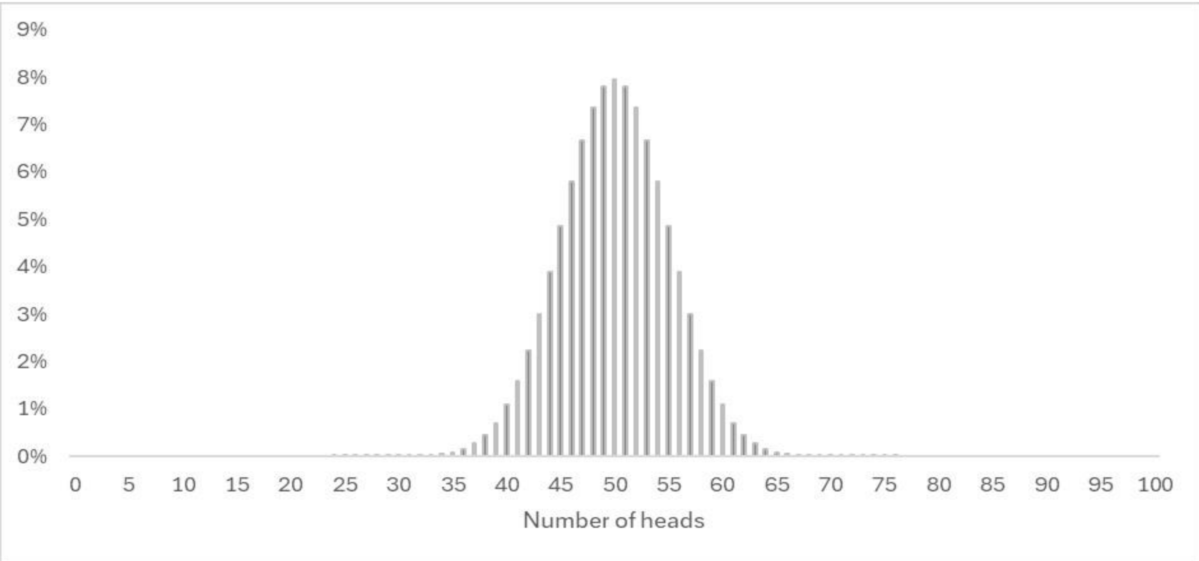

the other hand, predicts that every possible outcome of 100 coin flips should be surprising – that we ought to be surprised no matter what happens. The differential stance, finally, predicts that we should be surprised by some outcomes, like a run of 100 heads, but not others – which perhaps best accords with our pretheoretic intuitions. Focussing on surprise offers an interesting alternative perspective on these issues, but I won't explore this further here.

While looking at this distribution it seems plausible that I could justifiably believe that the number of heads will be somewhere *around* 50. More than 90% of the sequences, for instance, involve between 42 and 58 heads (inclusive) – so couldn't I justifiably believe that the number of heads will fall within this range? On the austere stance, the answer must be 'no'. After all, if I don't even have justification for dismissing the 'extreme' outcomes in which we have 0 heads and 100 heads then I couldn't be justified in believing that the number of heads will be around 50. I am assuming here that justification satisfies a single premise closure principle; if one has justification for believing a proposition P and P entails Q then one also has justification for believing Q. Given this principle, if I don't have justification for believing that there won't be 100 heads then I can't have justification for believing that the number of heads will fall within any non-trivial interval centered on 50. While I could still justifiably believe probabilistic claims like 'It's very likely that there will be between 42 and 58 heads', 'It's almost certain that we won't get 100 heads' and so on, I can't justifiably believe any out-and-out claims about the number of heads that will come up – except that it will be somewhere between 0 and 100 (inclusive). Thus the 'austere' stance seems well named.

Some might regard this result as unacceptably sceptical – if I'm asked to predict how many heads will come up, surely I could do better than just shrug and state that it will be somewhere between 0 and 100. And there is even a threat that this verdict could spill over into a more widespread scepticism. As is often observed, the number of heads resulting from a series of coin flips can provide a good, approximate model for the outputs of measurement processes that are subject to random noise. As a result, if I don't have justification for believing anything about the number of heads that will come up, it may be that a reading from, say, a thermometer or weighing scale would be incapable of providing me with justification for believing anything about the temperature or about an object's weight. I will return to this issue in section 5.

What, then, about the lavish stance? If underwritten by a probabilist theory of justification, the lavish stance will predict that I have justification for believing a lot of things about the number of heads that will come up – in fact, if anything, I will have justification for believing rather *too much* about this. Suppose we set the threshold for justification at 0.9. It will then turn out, as desired, that I have justification for believing that there will be between 42 and 58 heads. But I will also have justification for believing that there will be between 44 and 100 heads, that there will be between 0 and 56 heads and that there will *not* be exactly 50 heads – indeed I will have justification for believing any proposition that excludes less than 10% of the above distribution. While it might seem sensible to believe that there will be between 42 and 58 heads, believing that there will be between 44 and 100 heads or that there will be between 0 and 56 heads etc. looks quite arbitrary. And yet, on the probabilist theory, there is no reason why I should believe the first proposition in preference to any of the subsequent ones – they will all be equally worthy of belief. Worse still, there is nothing in the probabilist theory to prevent me from justifiably believing *all* of these propositions at once. If I'm asked how many heads will come up, surely it would be better for me to shrug than to reel off this motley of predictions.

What we want, perhaps, is the verdict that I have justification for believing that the number of heads will fall within an interval centered on 50 – say 42 to 58 – but *lack* justification for believing these myriad further propositions about the number of heads that will come up. What we need, then, is a kind of *differential* stance, on which I have justification for dismissing those sequences with an extreme number of heads – including 100 – but lack justification for dismissing the sequences in which the number of heads is closer to the average. This stance fits with the theory of justification developed by Goodman and Salow (2018, section 5, 2023, sections 4, 7).

As noted in the last section, Goodman and Salow deploy a formalism similar to that of the normic theory – in that it involves a set of possible worlds ordered with respect to their normalcy. On Goodman and Salow's theory, like the normic theory, the normalcy ordering yields a privileged set of worlds, such that one has justification for believing all and only the propositions that are true throughout that set. On the normic theory, this is the set of *maximally* normal worlds that are consistent with one's evidence. Goodman and Salow, however, allow for some gradations of normalcy within the privileged set, provided none of its members counts as 'sufficiently' less normal than any other.

This formalism might be applied to the coin flipping example as follows. Suppose we abstract away from all contingent features of the world other than the way in which the 100 coins land. This leaves us with 2^{100} worlds – one corresponding to each possible sequence that could come up. The most normal worlds will be those that feature exactly 50 heads. Those that feature 51 or 49 heads will be slightly less normal, those with 52 or 48 heads will be slightly less normal again and so on. Once we get to a very large or very small number of heads these small gradations will have added up to something 'significant' and the worlds will lie outside of the privileged set. Goodman and Salow define the 'sufficiently more normal than' relation in terms of both normalcy and probability: w_1 is sufficiently more normal than w_2 just in case the probability that things are more normal than w_2, given that they are no more normal than w_1, exceeds a threshold t (Goodman and Salow, 2023, pp121-122).

Assuming each world has a probability of $1/2^{100}$, if we set the threshold at 0.9 then the privileged set will include all and only those worlds that feature between 42 and 58 heads, and I will have justification for believing that the number of heads will fall within this interval. Goodman and Salow's theory ensures that there is a probability requirement for justification – any proposition I have justification for believing must have a probability higher than the threshold.[4] But, unlike the probabilist view, a high probability will not be enough for justification. So I won't have justification for believing that there will be between 44 and 100 heads, as there will be worlds in the privileged set that feature 42 and 43 heads. And I won't have justification for believing that there will be between 0 and 56 heads as there will be worlds in the privileged set that feature 57 and 58 heads. In fact, the proposition that there will be between 42 and 58 heads will, on this picture, be the *logically strongest* proposition about the outcome of the coin flips that I have justification for believing – that is, the only other propositions about the outcome that I have justification for believing will be propositions that are entailed by this one.

It's worth emphasising again that this model cannot be squared with the explanatory conception of normalcy from the last section. A fair coin landing heads requires no more explanation than a fair coin landing tails and is not the kind of event, on the explanatory conception, that could make a world any less normal. So what *is* normalcy then? Goodman and Salow, as I understand it, offer two possible answers. The first is to leave the notion of normalcy as a placeholder, and let our judgments about normalcy be guided by our epistemic intuitions (Goodman and Salow, 2023, p98) – including intuitive judgments about what we do and don't have justification for believing. This, in a way, shunts us over to a *methodological* question; are we just building a model to fit – and perhaps

---

[4] Let P be a proposition that holds throughout the privileged set. Let w_1 be a maximally normal world and let w_2 be a maximally normal world at which P is false. By the definition of 'sufficiently more normal' the proposition that things are more normal than w_2 given that they are no more normal than w_1 (which is trivial) must be greater than 0.9. Since the proposition that things are more normal than w_2 entails proposition P it follows that the probability of P must be greater than 0.9 as required. (This proof makes use of the simplifying assumption that, for any possible proposition, there are maximally normal worlds at which it holds – but can be adapted to do without it).

systematise – a set of intuitive judgments, or do we want our model to be independently motivated? Whatever we think in general about the model-fitting approach, it runs up against its limits in cases where our intuitions are *conflicted* – as they arguably are with the three stances. On this methodology, we would only have reason to accept Goodman and Salow's model if we were already convinced that the differential stance is the right way to go. But, as I will argue in the next section, I don't think we *should* be convinced of that.

In any case, Goodman and Salow also provide a more direct answer to the question of normalcy – which takes the form of a substantive account of what makes one state of affairs more normal than another (Goodman and Salow, 2021, section 2, 2023, section 5). One consequence of this account is that normalcy – and, by extension, justification – turns out to be something that is relative to a *partition*. That is, whether one possible world counts as more normal than another depends on how the space of possible worlds is divided up. I will turn to this in section 4.

3 Against the Differential Stance

I will raise two problems for the differential stance. Suppose I buy a ticket in an enormous lottery, and am assured by the organisers that the winner will be selected by a random process in which every ticket has an equal chance of winning. Suppose I then discover that the process for selecting a winner involves flipping 100 fair coins, with each possible sequence of heads and tails matched to one of the tickets. If the differential stance is correct then there will be some tickets, but not others, which are linked to sequences that I have justification for believing will not occur. Should I then complain to the organisers that the lottery is unfair, and insist that it be decided in a different way? I think it's pretty clear that I'm not in any position to complain. It's not as though the lottery organisers have misled me about anything – the winner *is* chosen by a random process and every ticket *does* have an equal chance of winning.

One of the supposed selling points of the differential stance is that it can deliver the result that I have justification for believing that the number of heads in 100 coin flips will fall within an interval around the average – say 42 to 58 – and that this is the logically strongest proposition about the outcome that I have justification for believing. Suppose I believe this proposition, and I then discover that my lottery ticket has been linked to a sequence with *59* heads, while my friend Bruce's ticket has been linked to a sequence with 58. Given my beliefs, my ticket is going to lose while Bruce's might still win. If I care about winning, wouldn't it be rational for me to try and *exchange* tickets with Bruce – perhaps even pay him for the privilege? In the original coin-flipping example, the question of how many heads will come up seems idle – whether I have any beliefs about this makes no obvious practical difference. But, as the present example demonstrates, these beliefs could certainly spur me into action if the setting is right and, in this case, would lead to actions that seem quite irrational.[5]

Consider another lottery example, in which the winning ticket is chosen by drawing numbered balls from a barrel. Suppose I discover that my ticket is linked to the sequence '535578' and, because of a superstition I have about the number 5, I come to believe that my ticket will lose. If I genuinely held this belief I would obviously be motivated to try and exchange tickets, and the present case is surely no different – if I genuinely believe that my ticket is going to lose once I discover that it's linked to a sequence with 59 heads, then I would try to swap if I can. And yet, these actions would seem

[5] This objection is inspired by the familiar point that one cannot rationally act on the proposition that one's ticket has lost a large fair lottery, prior to hearing about the result (see for instance Hawthorne, 2003, pp29-30).

equally irrational in both cases. In fact, when we compare the two beliefs back to back, it's hard to see why the latter belief shouldn't be regarded as *just as superstitious* as the former. And yet, according to the differential stance, this belief could be *justified*.[6]

That's the first problem for the differential stance – and the second is related. According to this stance, when I learn that 100 coin flips are about to happen I have justification for dismissing a sequence of 100 heads, but I don't have justification for dismissing the sequence HTTHTHTTTH…. Suppose I then discover that a yellow dot has been painted on one side of each coin – sometimes it happens to be the heads side and sometimes the tails side. This opens up the possibility of new remarkable outcomes – like the coins all landing with the yellow dots showing. But it's not as though the *total number* of possible outcomes has increased – a sequence of 100 yellow dots will simply correspond to some existing sequence of heads and tails. Suppose it corresponds to the sequence HTTHTHTTTH… What I have discovered, in effect, is a new way of *describing* this sequence, which makes it seem much more remarkable than it did before. By the same token, I also have a new way of describing the sequence of 100 heads that makes it seem much *less* remarkable than it did before – in terms of yellow dots and blank sides this will correspond to the sequence YBBYBYBBBY….[7]

If, instead of counting up the number of heads, we group the possible outcomes by counting up the number of yellow dots, we get exactly the same distribution – that is, an approximate normal distribution with a mean of 50 and standard deviation of 5. Once again, one of the alleged advantages of the differential stance is that it allows me to have justification for believing that the number of heads will fall within a certain interval around the average – 42 to 58 – while ensuring that this is the logically strongest proposition about the outcome that I have justification for believing. But once I'm made aware of the *yellow dots*, surely I have just as much justification for believing that there will be between 42 and 58 yellow dots as I do for believing that there will be between 42 and 58 heads. This leaves us with two alternatives – either the discovery makes me *lose* justification for believing that there will be between 42 and 58 heads or the discovery makes me *gain* justification for believing that there will be between 42 and 58 yellow dots.

---

[6] While a superstition about the number 5 might seem idiosyncratic, it's easier to imagine someone reacting negatively when they discover their lottery ticket is connected to a sequence with 'too many' heads. But that's not to say that this reaction would be any more rational. As I hinted in section 1, it's well known that people have a general tendency to associate randomness with a high level of alternation, and to dismiss the possibility of random processes producing outcomes that are 'patterned' or have a lot of repetition. But this is widely – and I think correctly – regarded as a *bias*, and not something that has a rational basis (see for instance, Falk, 1981, Lopes and Oden, 1987, Bar-Hillel and Wagenaar, 1991, Smith, 2016, section 3.4). I would be inclined to take things a step further, and argue that, in a way, this kind of tendency *couldn't* have any rational basis. When it comes to a random process, there is a sense in which the extent to which an outcome is 'patterned' is not even a feature of the outcome *itself* – but, rather, an artefact of the way in which we *choose to describe it*. The coming discussion might lend some support to this.

[7] As well as altering the coins in such a way as to introduce new remarkable outcomes, we could also take things in the opposite direction… Suppose the first coin has a heads side and a tails side, the second coin has a red side and a blue side (which are otherwise indistinguishable), the third coin has a '2' on one side and a '7' on the other and so on. In this case, every possible outcome will correspond to a specifiable sequence, and there will still be 2^{100} of them – but none will strike us as having a remarkable pattern, as there is no natural way to coordinate or match up the results of the individual coin flips (for a case like this see Bacon, 2014, pp391-392). When it comes to this case, a defender of the differential stance would be forced to adopt either a lavish stance or an austere stance, and would then be under some pressure to extend this to the original case as well. I think there are the makings here of another potential argument against the differential stance – related to, but distinct from, the argument pursued in the main text.

The first alternative does not seem promising. It's not as though the yellow dots affect the way the coins fly through the air or anything like that. Indeed, we can stipulate that the dots have no effect whatsoever on the coin flipping process. In that case it would be very strange if simply learning about the *existence* of the dots ended up defeating my justification for believing that there will be between 42 and 58 heads. While I could still perhaps retain justification for believing that the number of heads will fall within some broader interval, that brings little comfort. Whether there are yellow dots on the coins is just *irrelevant* to the question of how many heads will come up.

What about the second alternative? The same kind of issue arises. It's important to emphasise that, on this alternative, when I learn about the yellow dots I really will *acquire* justification for believing something about the outcome of the coin flips that I did not have justification for believing previously. The proposition that there will be between 42 and 58 yellow dots is equivalent to a proposition described exclusively in terms of heads-and-tails sequences – and this is *not* a proposition that I would have had justification for believing prior to discovering the yellow dots, as it's not entailed by the proposition that there will be between 42 and 58 heads. Just as it's strange to think that the discovery of the yellow dots should *weaken* my epistemic position with respect to how the coins will land, it's equally strange to think that the discovery of the yellow dots should *strengthen* my epistemic position with respect to how the coins will land.[8]

And even if we do embrace the second alternative, and insist that discovering the yellow dots really does give me justification for believing new things about how the coins will land, what happens if I then make further discoveries of this kind? Suppose I discover that one side of each coin has been marked with a green dot, some of which happen to be on the heads side, some on the tails side, some on the same side as the yellow dot, some on the opposite side… With each discovery of this kind, we face the same choice; either I lose justification for believing that there will be between 42 and 58 heads (and for believing that there will be between 42 and 58 yellow dots…) or I gain justification for believing another new thing – that there will be between 42 and 58 green dots etc. But after enough of these discoveries these propositions could even end up being *inconsistent* – that is, no matter how the coins land, there can't be between 42 and 58 heads *and* between 42 and 58 yellow dots *and* between 42 and 58 green dots *and* between 42 and 58 pink dots… and so on.[9]

[8] Another possibility is that the discovery of the yellow dots will both weaken *and* strengthen my epistemic position with respect to how the coins will land. That is, I could lose justification for believing some propositions about the outcome and gain justification for believing others. I don't just note this for completeness – as I will discuss in the next section, Goodman and Salow's theory, if developed in one natural way, will arguably lead to a prediction of this kind. In any case, this position won't resolve the problems identified in the main text – and may even exacerbate them. If it's strange to think that the discovery would weaken my epistemic position and strange to think that the discovery would strengthen by epistemic position, it may be doubly strange to think that it would have both of these effects.

[9] To illustrate the general phenomenon, suppose we have just four coins. Suppose there's a yellow dot on the heads side of the first two coins and on the tails side of the second two. Suppose there is a green dot on the heads side of the first and third coins and on the tails side of the second and fourth. Suppose there is a pink dot on the heads side of the first and fourth coins and on the tails side of the second and third:

	Coin 1	Coin 2	Coin 3	Coin 4
Heads side	YGP	Y	G	P
Tails side	-	GP	YP	YG

When we flip the coins the average number of heads/yellow dots/green dots/pink dots to come up will all be two – but, as can be easily checked, there is *no* possible outcome in which we have two heads, two yellow dots, The lavish and austere stances don't confront the two problems that I've identified in this section. In the coin-based lottery example neither stance will drive an epistemic wedge between my ticket and Bruce's. On the lavish stance I will have justification for believing, of each ticket, that it will lose, while on the austere stance I won't have justification for believing this about any tickets. In neither case will I have any incentive to swap. More generally, when I discover that the lottery is decided by 100 coin flips, according to the lavish and austere stances this makes no difference whatsoever to what I can justifiably believe about the outcome. Similarly, when I learn about the yellow dots, the lavish and austere stances predict that this makes no difference to what I can justifiably believe about how the coins will land. On the austere stance I don't have justification for believing that there will be between 42 and 58 yellow dots (or justification for believing that there will be between 42 and 58 heads). On the lavish stance, I *do* have justification for believing that there will be between 42 and 58 yellow dots – but this corresponds to a proposition which, when described purely in terms of heads-and-tails sequences, I would *already* have had justification for believing.

4 Partition Relativity

One thing that is distinctive about the yellow dots example is that it presents us with two natural ways of dividing up the possible outcomes – either by counting up the number of heads or by counting up the number of yellow dots. When we divide things up in the first way, it seems plausible that I have justification for believing that there will be between 42 and 58 heads, and when we divide things up the second way, it seems plausible that I have justification for believing that there will be between 42 and 58 yellow dots. We have been assuming, thus far, that justification is a binary relation between a subject and a proposition, but one reaction to this kind of example is to posit an *additional place* in the justification relation – a place to be filled by a *partition*, or a way of dividing the space of possibilities.

This would open up the following possible view: Relative to the number-of-heads partition I have justification for believing that there will be between 42 and 58 heads, but I lack justification for believing anything about the number of yellow dots. Relative to the number-of-yellow-dots partition I have justification for believing that there will be between 42 and 58 yellow dots but I lack justification for believing anything about the number of heads. One way to realise this view is by using Goodman and Salow's formalism, as explained in section 2.

Suppose again that we have 2^{100} worlds, representing the possible outcomes of the 100 coin flips. A partition of this set might be defined as any way of arranging the worlds into subsets – or cells – such that every world falls within exactly one cell. Given that the worlds are equiprobable, the probability of any given cell will be determined by dividing the number of worlds inside it by the total number (2^{100}). On the view Goodman and Salow propose, world w_1 is more normal than world w_2, relative to a given partition, just in case the cell that contains w_1 is more probable than the cell that contains w_2 (Goodman and Salow, 2021, 2023, section 5).[10] One potential advantage of this analysis

two green dots and two pink dots. In fact, every sequence that is average (two) for heads will be extreme (zero or four) with respect to one of the dot colours.

[10] What if we are dividing up the worlds according to the value of a *continuous* variable, such that there are an infinite number of cells, each with a probability of 0? In this case Goodman and Salow suggest that comparative normalcy be defined in terms of disparities in probability *density* (Goodman and Salow, 2023, pp120-121). (Think about the cumulative probability distribution associated with the variable – the probability that the variable will take a value no greater than x, for each x. The probability density, at the cell associated with value n, can be thought of as the rate of change, at n, of this cumulative probability). The number of heads/yellow dots in 100

is that it uses only probabilistic resources – resources that are available even to a probabilist. The price of this is to make normalcy into something that depends upon a partition – but, if the present considerations are on the right track, perhaps this could even be spun as an advantage as well. Recall that the 'sufficiently more normal than' relation, which determines the members of the privileged set, is defined as follows: w_1 is sufficiently more normal than w_2 just in case the probability that things are more normal than w_2, given that they are no more normal than w_1, exceeds a threshold t.

If we partition the worlds according to the number of heads then the maximally normal worlds will be those that involve exactly 50 heads – as this will be the largest cell in the partition. Those worlds that involve 51 or 49 heads will rank next in terms of normalcy as these will be the next largest cells and so on. If we set the threshold t at 0.9 then the privileged set will consist of all the worlds in which the coins are showing between 42 and 58 heads. These worlds are included in the largest cells in the partition and, between them, account for more than 90% of the total worlds. When we partition the worlds according to the number of yellow dots, the number and size of the cells remains the same – but the worlds themselves are shuffled around. For instance, the world in which there are 100 heads (= YBBYBYBBBY…) moves from a cell on its own into one of the much larger cells, effectively swapping places with the world in which there are 100 yellow dots (= HTTHTHTTTH…).

This will give us the desired upshot: Relative to the first partition I have justification for believing that the number of heads will be between 42 and 58, but I lack justification for believing anything about the number of yellow dots. Relative to the second partition I have justification for believing that there will be between 42 and 58 yellow dots but I lack justification for believing anything about the number of heads. Another way to partition the worlds is by counting up *both* the number of heads *and* the number of yellow dots that they involve. This partition will have a cell for 50 heads and 50 yellow dots, for 50 heads and 48 yellow dots and so on.[11] If we stick with a threshold of 0.9 then Goodman and Salow's theory will predict that, relative to this partition, I don't have justification for believing that the number of heads will be between 42 and 58 or for believing that the number of yellow dots will be between 42 and 58. If we assume that half of the yellow dots are found on a heads side and half on a tails side, then I will have justification for believing that the number of heads/yellow dots will fall within a certain broader interval – 40 to 60(inclusive).[12] Finally, we might consider a maximally fine-grained partition in which every one of the 2^{100} worlds is assigned its own cell. Relative to this partition, Goodman and Salow's theory predicts that I don't have justification for believing any non-trivial proposition about the outcome.

The partition-relative view can be thought of as a version of the differential stance – because, relative to a given partition, I may have justification for dismissing certain outcomes but not others. But it also enables us, in a sense, to be even-handed in the way that we treat the different outcomes – because, for any given outcome, there will be some partitions relative to which it can be justifiably dismissed and some partitions relative to which it cannot. Nevertheless, the idea that justification is

---

coin throws is, of course, a discrete variable – so we don't need this extra complication. In section 5 I will be considering the measurement of continuous variables such as weight, but even then the set of possible *measurements* will remain discrete, as measurement is not perfectly precise.

[11] Why not 50 heads and 49 yellow dots? You can't add or remove a yellow dot from a sequence without also changing a head to a tail or a tail to a head. Thus, if there are sequences that feature 50 heads and 50 yellow dots, there cannot be any sequences that feature 50 heads and 49 yellow dots.

[12] I omit the calculation here. The crucial point, in any case, is that there will be a broadening of the interval and, to see this, it suffices to point out that the proposition that there will be between 42 and 58 heads and between 42 and 58 yellow dots will, given the assumption in the main text, be less than 0.9, in which case it cannot hold throughout the privileged set of worlds.

partition-relative is one that carries significant costs – and it's unclear, on close inspection, whether it is able to help with the problems identified in the last section.

One thing we can immediately observe is that justification attributions have only two argument places at the level of surface form – one for a subject and one for a proposition. If justification is, in reality, a *ternary* relation linking a subject, a proposition and a partition, then the missing ingredient will need to come from somewhere. In Goodman and Salow's view, a partition will be supplied by the context in which a justification attribution is made and, in particular, will be determined by whatever *question* is salient in that context (Goodman and Salow, 2023, section 5).

If Goodman and Salow's theory is to give the desired result that I can truly assert 'I have justification for believing that the number of heads will be between 42 and 58' when I first hear about the coin flips, then the relevant partition, at this point, will have to be the one that sorts the worlds according to the number of heads that they involve – the partition corresponding to the question 'how many heads will there be?' If I then discover that one side of each coin has been marked with a yellow dot, it's plausible that the relevant partition will shift to one that sorts the worlds according to the number of heads *and* the number of yellow dots that they involve. If it helps, we could imagine someone asking the question 'How many heads and how many yellow dots will there be?' Relative to this partition, I have just as much justification for believing that there will be between 42 and 58 yellow dots as I do for believing that there will be between 42 and 58 heads. This leaves us with two options – either I can no longer truly assert 'I have justification for believing that there will be between 42 and 58 heads' or I can truly assert 'I have justification for believing that there will be between 42 and 58 yellow dots'. For the reasons noted above, Goodman and Salow's theory will favour the first option – relative to the new partition I would at most have justification for believing that the number of heads will fall between 40 and 60.[13]

On one level, these predictions are very similar to those of the *simple* differential stance. On the partition-relative view there is a sense in which the discovery of the yellow dots makes no difference to what I have justification for believing about the outcome of the coin flips; given a fixed partition, this doesn't change. But the discovery *does* make a difference to what I can truly *describe* myself as having justification for believing about the outcome of the coin flips. And – for me at any rate – it's hard to accept that the discovery should have even *this* effect. The fact remains that the presence of the yellow dots is simply irrelevant to the question of how the coins will land. When it comes to this kind of example, I don't see that the move to partition-relativity offers clear benefits to a defender of the differential stance.

The partition-relative view also places considerable strain on the relation between justification and rational *action*.[14] In the lottery example from the last section, I learn that my ticket has been linked to a sequence with 59 heads while my friend Bruce's ticket has been linked to a sequence with 58. In discussing the example, I endorsed the following conditional: If I have justification for believing that my ticket will lose and I lack justification for believing that Bruce's ticket will lose then it would be rational for me to try and exchange tickets with Bruce. We may want to shore this up by adding some further qualifications to the antecedent – that I do indeed want to win the lottery, that there are no independent reasons against trying to exchange tickets etc. – but, as long as we think there is some connection between justified belief and rational action, some version

[13] … and justification for believing that the number of yellow dots will fall between 40 and 60. Thus, the discovery of the yellow dots shifts us to a partition on which my epistemic position, regarding the outcome of the coin flips, will be weaker is some respects and stronger in others (see n8).

[14] The general thrust of this argument is similar to Hawthorne (2003, section 2.4).

of this conditional must be right.[15] Since the simple differential stance predicts that I *do* have justification for believing that my ticket will lose and lack justification for believing that Bruce's ticket will lose, the defender of this view faces a dilemma; either they have to reject the conditional, and sever the connection between justified belief and rational action, or they have to concede that it really would be rational for me to try and swap tickets. Neither option seems attractive.

What about the partition-relative view? On this view, there will be no 'absolute' fact about whether the antecedent of the above conditional is true. Relative to some partitions – like the number-of-heads partition – I do have justification for believing that my ticket will lose and I lack justification for believing that Bruce's ticket will lose. Relative to other partitions – like the maximally fine-grained partition – I don't have justification for believing, of any ticket, that it will lose. But the truth of the *consequent* does not depend upon a partition in this way. Suppose I do act in the ways described – when I discover the sequence linked to my ticket, I immediately complain to the organisers and pester Bruce for a swap. Either this is a rational reaction or it isn't, but one thing we can't say, it seems to me, is that my actions are rational relative to one partition while being irrational relative to another.

So the defender of the partition-relative view can only go one of two ways. If they deny that it would be rational for me to try and swap tickets then there will be some partitions relative to which the whole conditional is false. Relative to the number-of-heads partition for instance the conditional will have a true antecedent and a false consequent. If the defender wants to ensure that the conditional is true relative to all partitions then they will have to accept that the consequent is true – i.e. that it really is rational for me to try and swap tickets. When all is said and done, the defender of the partition-relative differential stance faces the *same dilemma* as the defender of the simple differential stance. Partition-relativity offers no way out.[16]

---

[15] Even with these qualifications, I don't think that this principle will necessarily carry over to any lottery scenario that we might imagine. A referee suggests the following example: Suppose I hold a ticket in a *two* ticket lottery, while Bruce holds a ticket in a billion ticket lottery with an equivalent prize. If one of the organisers tells me that the two ticket lottery is rigged in favour of the other ticket, then I would be justified in believing that my ticket will lose, while lacking justification for believing that Bruce's ticket will lose. But it is at best unclear whether it would be rational for me to seek a swap. This illustrates a more general point – when the stakes are high it may be irrational to take crucial propositions for granted, even if one does have justification for believing them (Smith, 2016, p86, section 5.2). The high stakes in this example might be explained in terms of the large discrepancy in the initial expected values of the two tickets. In any case, these kinds of considerations don't apply in the example described in the main text.

[16] In fact, matters may be even worse for the defender of the partition-relative view. Consider the following conditional, which seems just as plausible as the conditional above: If I have justification for believing that Bruce's ticket will lose and I lack justification for believing that my ticket will lose, then (given appropriate further qualifications) it would *not* be rational for me to seek to exchange tickets with Bruce. If we are attempting to be even-handed in the way we treat the different outcomes, then there should be some partitions relative to which the antecedent of this conditional is true. Goodman and Salow's theory will indeed predict that there are. Consider for instance the partition that sorts the outcomes according to the number of times they match the sequence associated with Bruce's ticket. That's a possible way of grouping the outcomes and, relative to this partition, I will have justification for believing that Bruce's ticket won't win, while lacking justification for believing that my ticket won't win (assuming the sequence associated with my ticket and the sequence associated with Bruce's have only an average number of matches).

We now have *two* plausible conditionals linking justification and rational action. If facts about justification are partition relative, but facts about rational action are not, then there must be partitions relative to which *at least one* of these conditionals is false. If the consequent of the first conditional is true then the consequent of the second conditional is false, and there will be partitions relative to which the whole conditional

5 A Problem for the Austere Stance

So far I have outlined a number of objections to the differential stance – in both its simple and partition-relative guises. I mentioned that my sympathies lie with the austere stance and, while I have noted some problems that might arise for this stance, I haven't subjected it to nearly the same level of scrutiny. In the final section I want to go some way towards redressing the balance. In section 2 I mentioned a worry that the austere stance could lead to a widespread scepticism concerning matters that are far removed from the outcomes of coin flips. Suppose I am about to measure the weight of an object using a sensitive digital scale. While I may not be justified in believing that the reading on the scale will precisely correspond to the weight of the object, I would surely be justified in believing that the reading will be somewhere around the true weight – that it will fall within a narrow interval centered on the true weight. And yet, this seems rather similar to the belief that the number of heads in 100 fair coin flips will be somewhere around 50 – a belief which, according to the austere stance, is unjustified.

The analogy between these two beliefs runs deeper than just initial impressions. There are any number of factors that could interfere with the reading on the scale – specs of dust, slight air movements, microphysical processes within the internal mechanism or electrical circuits etc. While any one of these factors would, at most, cause a minute divergence in the reading, in combination they could in principle lead to a larger error – though they are extremely unlikely to do so. If we were to plot the probability of the different readings we might obtain from weighing an object then, given plausible assumptions, this will approximate a normal distribution centered on the object's true weight. As we've seen, if we plot the probability of obtaining n heads in 100 coin flips, this will approximate a normal distribution centered on the mean number of heads – 50. Furthermore, the processes that can interfere with the scale reading are themselves often treated as being random – in the same sense as a coin flip – with their cumulative effects typically described as 'random noise'.[17]

If, as the austere stance contends, I don't have justification for believing anything about how far the number of heads in 100 coin flips will diverge from 50, and if a series of 100 coin flips provides a good model for the effects of random noise on the scale reading, it might seem to follow that I don't have justification for believing anything about how far the reading will diverge from the object's weight. If that's the case then, presumably, the reading on the scale couldn't provide me with justification for believing anything about how much the object weighs (Goodman and Salow, 2023, p133). But this conclusion would have us teetering on the brink of a very pervasive scepticism. A digital scale would be comparable, in these respects, to many measuring instruments – and even some of our own perceptual systems may be subject to a level of random noise.[18] Fortunately, I think the above reasoning can be resisted. After all, a *model* of a phenomenon – even a good one – is never a

---

is false. If the consequent of the second conditional is true then the consequent of the first conditional is false, and there will be partitions relative to which the whole conditional is false. There is no way for the defender of the partition-relative view to ensure the truth of both conditionals across all partitions.

[17] The approximate normal distribution of errors is standardly explained by citing the Central Limit Theorem which states, roughly, that the sum of a large number of independent random variables will be approximately normally distributed, irrespective of the distributions of the individual variables (see for instance Goodman, 2013, pp42-43). The theorem is nicely illustrated by the coin flipping case, in which individual coins can be construed as independent binary random variables that are summed to give the total number of heads.

[18] A fact which, according to Bricker (2019), has been largely overlooked by epistemologists and poses a threat to certain prominent epistemological views. While his primary target is the safety condition on knowledge, he implies that this fact is also uncomfortable for a normic theorist (Bricker, 2019, pp525-526). The coming material could I think be worked into a response to Bricker – but I can't give this the attention that it deserves here.

perfect fit with the real thing. In general, a good model will emulate certain features of a target phenomenon, while abstracting away from others.

Here is one respect in which weighing an object on a digital scale is not at all like flipping 100 coins: A scale is an instrument that performs a *function* – namely to measure the weights of objects placed upon it. The function of the scale is not to measure weight with perfect precision – it's clear that certain small errors in the reading will be compatible with the scale functioning properly. But it's equally clear that other, larger errors will *not* be compatible with the scale functioning properly. And this is so irrespective of what *explains* the error – whether it is miscalibration or deliberate tampering or the cumulative effects of random interference. If the error is too large, the scale is not functioning as it should. The boundary between the error magnitudes that are compatible with proper functioning and the error magnitudes that are not will, of course, be vague – that is, characterised by borderline cases. The boundary will also depend, in non-obvious ways, on the particular design of the scale. But none of this is to say that there *is* no boundary.[19] To insist that any error magnitude is compatible with the proper functioning of the scale is to give up on the idea that the function of the scale is to measure weight (and, arguably, to give up on the idea that deviations from the true weight should even be described as 'errors').

So weighing an object on a scale is a process that has *teleological* structure – and this structure is missing in the case of flipping 100 coins. A set of 100 coins doesn't have the 'function' of showing around 50 heads when flipped – and deviations from this number are not 'errors' in any sense. A series of coins that all land heads is functioning just as well as a series in which half the coins land heads and half land tails. Whether this teleological structure should have any significance for epistemic justification is a further question but, at the very least, there are resources here to which the defender of the austere stance can appeal in order to head off a full-blown sceptical result.

The normic theory of justification may not assign direct significance to teleological considerations – but it's clear that a teleological structure can, in some cases, drive apart the normalcy of possible outcomes.[20] If the scale produces a reading that diverges too far from the true weight then *something* has interfered with its proper functioning. Whether this is down to tampering or miscalibration or random processes, the divergent reading will be *explicable* and, as a result, the normic theory will predict that I have justification for believing that the reading will not diverge far from the true weight. If the possibility of random interference is compatible with the scale having a function then it is also compatible with some readings being more normal than others. When it comes to the possibility of random interference, the part to focus on is not the 'random' but the 'interference'.

So the defender of the normic theory, and of the austere stance more generally, can insist that there is a crucial disanalogy between a series of coin flips and weighing an object on a scale – a disanalogy between a random process and a goal-directed process that may be subject to random interference.[21] In a way, the trickiest problem for the austere stance arises in a case in which we have

[19] See the related discussion in Smith (2018a, pp732-734).

[20] See the related discussion in Silva (2023, pp2656-2657).

[21] In addition to the worry about measurement, Goodman and Salow raise another sceptical problem for the normic theory, focussing on its prediction that one cannot justifiably believe that a coin is double headed on the grounds that it has been observed to land heads many times in a row. But this, according to Goodman and Salow, is analogous to the way in which we believe in laws of nature, based upon observed regularities (Goodman and Salow, 2023, p133, see also Bacon, 2014, pp377-378). I disagree that this is a good model for how we believe in laws of nature – but this will have to be a topic for another occasion.

specific information about the random processes that might interfere with the scale reading – for information like this can trigger a shift in perspective, and *defeat* our justification for believing that the reading lies close to the true weight. The information in question is not just that scale readings are *sometimes* disrupted by random processes – this would be comparable to the information that scales are sometimes miscalibrated or sometimes subject to deliberate tampering, and will not be enough to have a defeating effect. Justification is compatible with a mere awareness of error possibilities – it's only when we have evidence in favour of such possibilities that justification is threatened. But when it comes to evidence that relates to random interference, it seems that relatively little is needed.

Suppose I discover a random process with a specified outcome that would result in a significant deviation in the scale reading. If I had justification for believing that the reading and the weight were close then, given my additional discovery, I would also have justification for believing that the random process in question did not result in that one specified outcome. And, according to the austere stance, this is not something that I could have justification for believing. This discovery is quite precise, in that it connects a significant deviation in the reading with one particular outcome – but the threat of defeat may extend further than this. If I know that some outcome of an unfolding random process would cause a significant deviation in the reading then, on the austere stance, this could still result in defeat, even if I don't know *which* outcome it is.

To put this in a schematic way, suppose I know that a random process R has played out and that one of its outcomes would have significantly disrupted the scale reading. Suppose o is one of the possible outcomes of R. For all I can justifiably believe, o might be the outcome that disrupts the reading and, for all I can justifiably believe, o might be the outcome that eventuated. If these two propositions are independent, given my evidence, then, for all I can justifiably believe, o might be the outcome that disrupts the reading *and* the outcome that eventuated. In this case, I could not justifiably believe that the reading lies close to the true weight.[22] While I can take a teleological stance towards the scale and the reading – and not towards a series of coin flips – this stance is, in one respect, quite fragile. If I have the right kind of information about random processes impinging on the reading then this could force me to see it in a different light, on which it can only provide justification for believing probabilistic claims about the weight of the object – 'it's very likely to be close to the reading' etc.

Whatever we make of this though, I don't think it can be a decisive consideration in the dispute between the austere and differential stances – because the differential stance is going to lead to a very similar result. To keep things as simple as possible, let's return to the case with which we began, and imagine that the 100 coin flips takes on the role of an 'interfering' random process. With apologies to Schrödinger, suppose I know that a cat, who was alive and well, was moments ago placed in an air-tight box. Suppose I discover that there is one specified sequence of heads and tails such that, if the 100 coins land in that sequence, the box will be filled with lethal gas, otherwise the cat will be released

---

[22] Suppose evidence E provides justification for believing P. D might be described as a defeater just in case $E \wedge D$ does not provide justification for believing P. Another way to think about the present issue is in terms of the following principle: If D_1 is a defeater ($E \wedge D_1$ does not provide justification for believing P) and D_2 is a defeater ($E \wedge D_2$ does not provide justification for believing P) then $D_1 \vee D_2$ is a defeater ($E \wedge (D_1 \vee D_2)$ does not provide justification for believing P). This is a 'defeat' version of the *amalgamation* principle discussed in Smith (2018, pp3866-3867, see also Smith, 2022b, n14) – and can of course be extended to any number of disjuncts. If E is the background evidence, P is the proposition that the reading lies close to the true weight and the D_is, are propositions linking each possible outcome of the random process to a significant disruption in the reading, then this principle gives us another route to the conclusion in the main text.

unharmed. Suppose the coins have been flipped but I haven't seen the result. After making this discovery, do I have justification for believing that the cat is alive? On the austere stance, the answer is 'no' – when I make the discovery, my justification for believing that the cat is alive is defeated. On the lavish stance the answer is 'yes'. On the lavish stance I already have justification for believing, of each sequence, that it won't come up, so this discovery doesn't really affect anything.

On the differential stance, the answer depends on what the key sequence *is*. If the sequence linked to the gas happens to be one of those that we already had justification for dismissing – like a run of 100 heads – then the discovery makes no difference, and I would have justification for believing that the cat is still alive. But if the sequence happens to be one of the others – like HTTHTHTTTH… – then the discovery *will* make a difference. In this case, I will either *gain* justification for believing that the coins did not land in the sequence HTTHTHTTTH…, which doesn't seem plausible, or I will *lose* justification for believing that the cat is alive, just like on the austere stance. Furthermore, if I discover that there is some sequence of heads and tails that will trigger the gas, without discovering which sequence it is, then, on the differential stance, this should also be enough for defeat to take place. If I know that the selection of the lethal sequence and the outcome of the coin flips are independent of one another then the differential stance would seem to predict that, for all I can justifiably believe, HTTHTHTTTH… might be the sequence that triggers the gas and also the sequence in which the coins landed. In that case, the differential stance predicts that I don't have justification for believing that the cat is still alive. So when it comes to this case, the simple differential stance doesn't seem to offer any better predictions than the austere stance.

Finally, consider Goodman and Salow's partition-relative view. As we've discussed, when we first hear about the coin flips, we want the relevant partition to be the one that divides the worlds according to the number of heads showing. It is relative to this partition that I have justification for believing that the coins won't all land heads and I lack justification for believing that the coins won't land in the sequence HTTHTHTTTH…. Relative to this partition, things will play out in the same way. That is, if I discover that the sequence HTTHTHTTTH… is linked to the death of the cat, or I simply discover that there is some sequence that is linked to the death of the cat, I would no longer be justified in believing that the cat is alive. If this is the relevant partition then it would be true for me to say 'I don't have justification for believing that the cat is alive' and we get a result very like the austere stance.

But the defender of the partition-relative view has another option of course – it could be that the discovery changes the relevant partition. Suppose I discover that the sequence HTTHTHTTTH… is linked to the death of the cat and this shifts us to a partition on which I *do* have justification for believing that the cat is alive and also have justification for believing that the coins did not land in the sequence HTTHTHTTTH…. It's unclear to me why the discovery should have that effect but, even if it does, we are still left with a troubling result. Before I make the discovery I could truly say 'For all I have justification for believing, the coins might have landed in the sequence HTTHTHTTTH…'. After I make the discovery I can truly say 'I have justification for believing that the coins did not land in the sequence HTTHTHTTTH…'. But surely the discovery can't put me in a position to truly make this claim – learning that a given sequence would lead to surprising downstream consequences doesn't give us reason for thinking that the sequence won't come up.[23]

[23] If I were to discover that the sequence HTTHTHTTTH… involved, say, 90 yellow dots then the defender of the partition-relative view is, of course, already committed to the claim that I could then truly say 'I have justification for believing that the coins didn't land in the sequence HTTHTHTTTH…'. So perhaps there's little additional cost in thinking that the lethal gas discovery could have the same effect. On the other hand, there just seems to be

In this paper I've set out an epistemic puzzle which is ostensibly about coin flips but, more generally, is about random processes that can produce striking or remarkable outcomes. I've identified three stances that one can adopt in response to the puzzle; the lavish stance which fits with the probabilist theory of justification, the austere stance which fits with the normic theory of justification that I've defended in previous work, and the differential stance which fits with the theory of justification recently defended by Goodman and Salow. While I'm sympathetic to the austere stance, and decided to organise the paper around that, this may of course be *because* it fits with views that I've previously defended, as much as its intrinsic merits. In truth, I don't see that the considerations I've offered here support any final conclusions as to which stance is right. Rather, the main contribution has been to show just how deep the puzzle goes and how much strain it places on our ordinary notion of justified belief.

Acknowledgements

Thanks to an anonymous referee for very helpful comments and suggestions. Work on this paper was supported by the Arts and Humanities Research Council (grant no. AH/T002638/1).

References

Bar-Hillel, M. and Wagenaar, W. (1991) 'The perception of randomness' *Advances in Applied Mathematics* v12(4), pp428-454

Bacon, A. (2014) 'Giving your knowledge half a chance' *Philosophical Studies* v171(2), pp373-397

Bricker, A. (2019) 'The myth of stochastic infallibilism' *Episteme* v18(4), pp523-538

Falk, R. (1981) 'The perception of randomness' *Proceedings of the Fifth International Conference for the Psychology of Mathematics Education* (Grenoble: Laboratoire IMAG)

Goodman, J. (2013) 'Inexact knowledge without improbable knowing' *Inquiry* v56(1), pp30-53

Goodman, J. and Salow, B. (2018) 'Taking a chance on KK' *Philosophical Studies* v175(1), pp183-196

Goodman, J. and Salow, B. (2021) 'Knowledge from probability' in Halpern, J. and Perea, A. eds. *Theoretical Aspects of Rationality and Knowledge* (Waterloo, Australia: Open Publishing Association)

Goodman, J. and Salow, B. (2023) 'Epistemology normalized' *Philosophical Review* v132(1), pp89-145

Hawthorne, J. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)

Lando, T. (2017) 'Coincidence and common cause' *Noûs* 51(1), pp132-151

Lopes, L. and Oden, G. (1987) 'Distinguishing between random and non-random events' *Journal of Experimental Psychology* v13(3), pp392-400

something back to front about this reasoning. If I learn that the sequence HTTHTHTTTH… would, surprisingly, lead to some consequence c, then that should change how I can truly describe my epistemic position with respect to *c*, but shouldn't change how I can truly describe my epistemic position with respect to HTTHTHTTTH….

Silva, P. (2023) 'Merely statistical evidence: When and why it justifies belief' *Philosophical Studies* v180(9), pp2639-2664

Smith, M. (2010) 'What else justification could be' *Noûs* v44(1), pp10-31

Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief* (Oxford: Oxford University Press)

Smith, M. (2017) 'Why throwing 92 heads in a row is not surprising' *Philosophers' Imprint* v17(21), pp1-8

Smith, M. (2018a) 'Coin trials' *Canadian Journal of Philosophy* v48(5), pp726-741

Smith, M. (2018b) 'The logic of epistemic justification' *Synthese* v195(9), pp3857-3875

Smith, M. (2022a) 'The hardest paradox for closure' *Erkenntnis* v87(4), pp2003-2028

Smith, M. (2022b) 'Two accounts of assertion' *Synthese* v200(3), pp1-18