King's Research Portal

DOI:

[10.1007/s11098-019-01276-2](https://doi.org/10.1007/s11098-019-01276-2)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Littlejohn, C., & Dutant, J. (2019). Justification, Knowledge, and Normality. *PHILOSOPHICAL STUDIES*.
<https://doi.org/10.1007/s11098-019-01276-2>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Justification, Knowledge, and Normality

Clayton Littlejohn

Julien Dutant

Forthcoming in Philosophical Studies. Please cite the published version.

There is much to like about the idea that justification should be understood in terms of normality or normic support (Smith 2016, Goodman and Salow 2018). The view does a nice job explaining why we should think that lottery beliefs differ in justificatory status from mundane perceptual or testimonial beliefs. And it seems to do that in a way that is friendly to a broadly internalist approach to justification. In spite of its attractions, we think that the normic support view faces two serious challenges. The first is that it delivers the wrong result in preface cases. These cases suggest that the view is either too sceptical or too externalist. The second is that the view struggles with certain kinds of Moorean absurdities. It turns out that these problems can easily be avoided. If we think of normality as a condition on *knowledge*, we can characterise justification in terms of its connection to knowledge and thereby avoid the difficulties discussed here. The resulting view does an equally good job explaining why we should think that our perceptual and testimonial beliefs are justified when lottery beliefs cannot be. Thus, it seems that little could be lost and much could be gained by revising the proposal and adopting a view on which it is knowledge, not justification, that depends directly upon normality.

1. Introduction

Recall Harman's observations concerning the lottery:

In the testimony case a person comes to know something when he is told about it by an eyewitness or when he reads about it in the newspaper. In the lottery case, a person fails to come to know he will lose a fair lottery, even though he reasons as follows: "Since there are N tickets, the probability of losing is $(N-1)/N$. This probability is very close to one. Therefore, I shall lose the lottery." A person can know in the testimony case but not in the lottery case ... The contrast between the two cases may seem paradoxical, since witnesses are sometimes mistaken and newspapers often print things that are false. For some N , the likelihood that a person will lose the lottery is higher than the likelihood that the witness has told the truth or that the newspaper is right (1968: 166).

When we first reflect on this, this does seem paradoxical. It is odd to think that it is easier to know the less probable proposition. Things go from strange to stranger when we think about justification. We can justifiably believe that Arsenal lost when *The Guardian* reports that, but we cannot justifiably believe that a ticket lost in a fair lottery, not even if the number of tickets sold is very large (Nelkin 2000). If our confidence ought to match the probabilities, that means that we ought to be less confident of the proposition that we are justified in believing. How could that be?

One recent proposal cuts these two lottery problems down to one. According to Sutton (2007), justified beliefs just *are* pieces of knowledge.¹ If so, the explanation as to why testimonial beliefs are justified and lottery beliefs are not just follows from our preferred explanation as to why the former can be known when the latter is not (e.g., only one kind of belief can be sensitive, safe, etc.). There is one lottery paradox, not two. Once we remember that the anti-luck conditions necessary for knowledge have little to do with the probabilities of the relevant propositions, the

problem with knowledge doesn't seem all that pressing. And on that view, once that problem has been solved, there is no further problem about justification to address.

The problem with this proposal, some might say, is that it seems too externalist.² Our impression is that a critical mass of epistemologists think something like this. On the one hand, it is right and proper for reliabilists to get their chance to defend their account of justified belief in the journals even if they might struggle with the intuitions that some of us have when we consider Cohen's (1984) new evil demon objection. On the other, it is right and proper for us to balk at the suggestion that justification is factive. These factive views are too externalist. Here is a stab at articulating the objection. Suppose some subject forms some perceptual or testimonial beliefs. Most of these beliefs are knowledge and all could be knowledge if things turn out well. Suppose that from the inside each belief appears to be a perfectly good candidate for knowledge. A theory that treats these beliefs as differing in justificatory status is *spotty*; a theory that treats them alike is *smooth*. If a theory is spotty in spite of how good things look from the inside, in spite of the absence of some discernible defeater, and in spite of how reliable the type of process is in the situation in which it is operative, then, the thought goes, it is too externalist.³

Here we consider an alternative account of justification that is designed to vindicate our intuitions about testimonial beliefs and lottery beliefs while avoiding the "too externalist" objection. Consider Smith's (2016) normic support view:⁴

Normality. A thinker has justification to believe p iff the thinker's evidence E provides normic support for p , that is, p is true at the most normal worlds in which E is true.⁵

Crucially, Smith's notion of normalcy is cashed out in terms of explanation rather than statistical frequency. The intuitive idea is that abnormal occurrences call for explanation, but normal ones do not. As Smith puts it:

In this sense of 'normal' it could be true that Tim is normally home by six, even if this occurrence is not particularly frequent. What is required is that exceptions to this generalisation are always explicable as exceptions by the citation of independent, interfering factors—his car broke down, he had a late meeting, he had to detour around roadworks etc (2016: 40).

If a ticket in a fair lottery wins, there is no need to explain why it did. However statistically unlikely, that is a normal occurrence. By contrast, if your vision misfires or your reliable witness gives false testimony, an explanation is called for: the light was tricky, there was a similar-looking person in the area, and so on. Thus believing without normic support for your belief puts you at risk of

² An anonymous referee raised the reasonable concern that describing views like Sutton's as 'too externalist' might needlessly annoy card-carrying externalists. The authors of this paper understand this concern. One of them *is* a card-carrying externalist, one who thinks that Sutton was right when he said that a belief is justified iff that belief constitutes knowledge. Because of this, this author is frequently told that his view, unlike reliabilist views, is 'too externalist' to merit serious consideration. We hope that readers won't be bothered by the casual use of this talk of a view being 'too externalist' when they know that one of the authors talking this way is frequently a target of this kind of criticism. For what it is worth, this author likes the accounts sketched in §4 because he sees the need for a more subjective, more internalist mode of epistemic evaluation. He prefers to think of this kind of evaluation as more concerned with rationality than with doxastic justification, but he is happy to go along with the idea that justification should be understood as something like rationality and something that depends primarily upon features of the thinker's mental life.

³ This objection is heard often in conversation. See Gerken (forthcoming) and Huemer (2006) for similar ideas in print. Note that on Huemer's view even standard reliabilists face the problem. Since the views proposed here are compatible with Huemer's idea that justification supervenes upon an individual's non-factive mental states, we need not worry too much about whether his charge against standard reliabilism is warranted.

⁴ See also Goodman and Salow (2018). We focus on Smith's more detailed account. In §4 we suggest endorsing Goodman and Salow's account of knowledge, not their account of justification.

⁵ More precisely: E normically supports p iff for any E -world at which p is false, there is a more normal E -world such that p holds at all E -worlds at least as normal as it. See Smith (2016, §7.1, §8.1). The simpler version will do for our purposes.

¹ Littlejohn (2012), McDowell (1998: 395), Praolini (forthcoming), and Williamson (2013) also defend views on which doxastic justification requires truth.

“brute or inexplicable error” in a way that believing with normic support does not (Smith 2016: 78). Like Sutton’s knowledge account, the view vindicates the intuition that you don’t have adequate justification to believe lottery propositions without implying that our testimonial beliefs must be unjustified by virtue of being less probable on the evidence. Unlike Sutton’s account, however, it allows for false, justified belief. It allows that something abnormal can happen where this abnormal occurrence accounts for the belief being false without making it unjustified. Thus, it seems to vindicate an important intuition often offered in support of a knowledge-first approach to justification without embracing the ideology. For many, this would be a selling point.

Let’s consider another attractive feature of the view. Consider this lottery-type example:

Prisoners. 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The 100th prisoner played no role in the assault and could have done nothing to stop it. There is no further information that we can use to settle the question of any particular prisoner’s involvement (Redmayne 2008).

Many of us think that it wouldn’t be appropriate to punish a prisoner chosen at random from the yard. Many of us think, however, that it might be appropriate to punish if an eyewitness came forward to testify that some particular prisoner was involved with the assault. Smith’s account vindicates both judgements. We don’t have normic support to believe that a randomly picked prisoner was involved. Given our evidence, if that prisoner turned out not to be involved with the assault that would not call for an explanation. For the conjunction of your evidence with the proposition that they are involved requires no more explanation than the conjunction of your evidence with the proposition that they are not. By contrast, the conjunction of an eyewitness’s testimony against a prisoner with the proposition that they are not involved more explanation than its conjunction with the proposition that they are. So the view predicts that an eyewitness’s testimony would justify believing that that prisoner was involved. Provided that it is appropriate to punish a prisoner if and only if it is justified to believe that they are involved, the view vindicates widely shared intuitions. Smith (2018) sees this as a virtue of Normality.⁶

Fans of the normic support view account for the difference between lottery beliefs and testimonial beliefs as a difference in whether the beliefs enjoy normic support. When cashing out the view in formal terms, they will say that a thinker’s evidence *E* provides normic support for *p* at a world *w* iff there exists a P-permitting normalcy sphere associated with *w* in which all the *E*-worlds are *p*-worlds (Smith 2016: 117). We are supposed to imagine spheres of worlds where the innermost spheres are, from the perspective of *w*, the most normal and to hold the evidence fixed and check whether the error possibilities are at least as normal (i.e., at least as close to the innermost sphere) as the possibilities in which the target proposition is true. We might naturally wonder whether a view that models the kind of evidential support supposed to be involved in the justification of belief in that way matches up with our intuitions about justification. Our fear is that the model cannot vindicate the intuitions that motivate it in the first place. Indeed, we think that the intuitions that seem to favour this approach actually clash with the normic support view and that the best way forward would be to reject it.

2. The Normic Preface

Let’s consider a simple example:

⁶ There are alternative proposals as to why it is not appropriate to punish on the basis of statistical evidence alone. In these debates, many of us want to follow Adler (2002) and Buchak (2013) in saying that it would not be proper to punish or hold someone responsible if it were not proper to believe that they were guilty. On Smith’s proposal, the lack of justification for the relevant beliefs about guilt is chalked up to the fact that the thinker does not have the right kind of evidential support. See Moss (2018) for the claim that further, non-epistemic factors are involved and Gardiner (forthcoming) for a response.

Eye Exam. An eye doctor has you look into a device that flashes slides. You were asked to identify the number or letter that you see. In each case, the letter or number appears quite clearly to you. The exercise is a bit tedious as it goes on for quite some time without much by way of variability. The letters and numbers aren’t particularly small or blurry. Your doctor says that you did very well and won’t need glasses.

This is a scenario where you would seem to acquire loads of perceptual knowledge. Plausibly, you would thereby acquire loads of justified perceptual beliefs. And if you had a good memory, you might also remember the series that you saw. Suppose, however, you press your doctor to tell you how well you did. She says that you did very well and that you made only one mistake.

Here is a natural thought. When it seems to you that *p* because of what you seem to perceive or what you are told, you are justified in believing *p* provided that there are no available defeaters. Such an account fits well with internalist and non-skeptical intuitions (Pollock 1974, Pryor 2000). If we adopted it, we would say that before your doctor tells you that you made a mistake, each of your perceptual beliefs is justified. On its own, the doctor’s testimony would itself justify a testimonial belief that you made one mistake. Furthermore, we might think that her testimony isn’t a defeater for any of your perceptual beliefs about a single letter or number. For her testimony does not tell you that *that* belief is false. (Nor does it have a significant impact on its probability or reveal that the beliefs initially lacked normic support.) It only tells you that *their* conjunction is false. Conversely, we may think that none of your individual perceptual beliefs defeats her testimony. For none of them tells you that you didn’t make any mistake. On that conservative take on defeat we would say that all of your individual perceptual beliefs are justified as well as the belief that one of them is mistaken. The set of justified beliefs would be inconsistent and it would not be closed under conjunction, but not everyone thinks that that’s a bad thing.⁷

Here is another natural thought. If a subject is internally like someone who knows *p*, it is justified for them to believe *p*.⁸ Now it seems that you could have acquired perceptual knowledge as you went through the eye exam and that you could have acquired testimonial knowledge in hearing from the doctor that one mistake was made. Of course since the contents in question form an inconsistent set you could not have known all together. But *each* content is something that an internally indiscernible thinker could have known. So on that thought as well, all of your perceptual beliefs about particular letters are justified, as well as the belief that one of them is mistaken. Again, your justified beliefs are neither consistent nor closed, but that may be defensible.

On two approaches to justified belief we have a plausible argument for the following:

Justified Inconsistency. After you are told by the doctor that you made one mistake, you are in a position to justifiably believe an inconsistent set of propositions about the letters on the cards (i.e., that the first card was a ‘P’, the second card was an ‘R’, the third card was an ‘E’, the fourth card was an ‘F’, the fifth was an ‘A’, etc. and that one of these beliefs is mistaken).

Justified Inconsistency is a problem for the normic support view. For on that view you are only justified in believing propositions that hold at all the most normal worlds at which your evidence

⁷ The preface was first introduced into the literature by Makinson (1965). For further support for the idea that there can be justified or rational inconsistency, see Christensen (2004), Easwaran and Fitelson (2015), Foley (1992), and Worsnip (2016). For dissenting views, see Adler (2002), Evnine (1999), Leitgeb (2017), Pollock (1986), Roush (2010), and Ryan (1991).

⁸ See Bird (2007) for a defense of this view. We shall consider similar proposals below. For our purposes, subjects are ‘internally alike’ when they are in the same non-factive mental states (e.g., the one seems to see what the other sees, seems to remember what the other remembers, believes what the other knows, etc.). On this way of thinking about things, you might have an internal duplicate who has recently been envatted but would not have an internal duplicate on Twin Earth if the thinkers located there have propositional attitudes with different contents. Thanks to an anonymous referee for raising this issue.

holds. Since the propositions you are justified in believing in hold at all most normal worlds, they hold at some world, and since they hold at some world, they must be mutually consistent.⁹ So the propositions referred to in Justified Inconsistency cannot all receive normic support from a thinker's evidence.

Once we see that the view must reject Justified Inconsistency, we should try to work out what its proponents would say about Eye Exam. At first they may suggest that it is an instance of the Preface paradox. For you have independent, non-statistical evidence for each of your perceptual beliefs, just as an author may have independent, non-statistical evidence for each of the claims in a book. Yet you have good reason to think that one of these beliefs is false, just as an author as good reason to think that their book contains a mistake. In his treatment of the Preface, Smith argues that the claim that the book contains a mistake is not justified. For it is normal, though perhaps unlikely, for each claim in the book to be true. By contrast, any error, wherever it is, would call for an explanation (2016: 74). However, in Smith's version of the Preface, the author merely has "strong inductive evidence" (e.g., statistical evidence) that their book contains an error. The only thing that the mistake claim has going on for it, so to say, is that it is highly likely to be true. If it was false no explanation would be called for. In fact, Smith sees the Lottery and the Preface paradoxes as mirror cases: in the former one has statistical evidence for the conjuncts but normic evidence for the negation of their conjunction while in the latter one has normic evidence for the conjuncts and statistical evidence for the negation of their conjunction (2016: 75). In both Smith maintains that one's justified beliefs are consistent by arguing that statistical support is not enough for justification.

The Eye Exam case is different, however. The support you have for the mistake claim is precisely of the kind that Smith wants to classify as normic rather than statistical. If the eye doctor's report turned out to be false that would require an explanation: that they didn't pay attention, they failed to read the letters correctly, or the like. So Smith's diagnosis of the original Preface does not carry over. The structural feature of the case is that your evidence makes it certain that something abnormal happened: either that you didn't correctly see a letter or that the doctor's report is incorrect. A newspaper headline once said something like, "local football team expects no surprises this season". That is newsworthy because it is sometimes reasonable for a team to expect surprises. For there may have a list of normal outcomes of each game and still say that it be normal for some game to go otherwise than expected. In conversation, a few people have expressed sympathy for the idea that there might be a series of events such that it would be abnormal for each not to occur, and yet abnormal for all to occur. The Eye Doctor case has that structure. What should proponents of the normic view say about such cases?

One option is to say that only some of the individual perceptual beliefs are justified. If there is some letter such that it would be more normal for your evidence to hold and you being mistaken about it than for your evidence to hold and you being mistaken about any other letter or number, the normic view can say that you are not justified in believing your answer on that letter. (Similarly, if we held that a false report for the eye doctor is more normal than any mistaken perceptual belief of yours, the view could say that it is not justified to believe the doctor's report, and perhaps even justified to believe that it is false.) That need not mean that any of these beliefs was unjustified before the doctor told you that you made a mistake. For it may still be that among the worlds in which the doctor merely told you that you did well and didn't need glasses the most normal ones are those in which you did not make any mistakes. Now that option is an interesting theoretical possibility. But it does suggest that there are secret features of the normality structure that pick winners and losers from the beliefs on the basis of factors that a thinker could not possibly discern. It is a spotty view. It is too externalist. It gives up the outlook that initially motivates the normic support view over Sutton's.

⁹ This assumes that one's evidence holds at some world in the actual world's system of spheres, hence at the very least that it is consistent. Pleading that one's evidence is inconsistent in the case would not help, however, as inconsistent evidence provides normic support for any proposition whatsoever.

A second option is to say that all your individual perceptual beliefs are defeated. The idea is that provided that in a scenario there is nothing from the inside to pick out some perceptual beliefs from the others, an error in one of the perceptual belief is just as normal as an error on any other. After the doctor's report the most normal worlds compatible with your evidence are ones in which one perceptual belief is an error. So your evidence fails to provide any of them with normic support. Unlike the Lottery, the Eye Doctor involves a *prima facie* support for each of the particular conjunct propositions that is not statistical in kind but normic-like. As in the Lottery, however, one's evidence makes it most normal that one conjunct is false but does not make it more normal for any one to fail rather than another. So on the normic view, there is no justification to believe any. (On a related view, mistakes on any of the particular perceptual beliefs are as normal as a mistake in the doctor's report. On that view one is not justified in believing the doctor's report either.) By preventing the normality structure from playing favourites among beliefs that look equally good from the inside, the view remains smooth.

That option is not as sceptical as it may first seem. For while it entails that individual conjuncts like "this card was an E" are not justified, it allows that more general propositions are justified. For instance, we may think that in the light of your evidence after the doctor's report it would be more normal for you have made *one* mistake rather than *two*. If so the normic view says that you're justified in believing that all but one of your answers were correct. That is to say, you're justified in believing the disjunction of every maximally consistent conjunction of target propositions. So we need not say that all your perceptual beliefs are defeated.

We fear, however, that the option is too sceptical still. We didn't choose Harman's quotation at random. It's baked into the lottery puzzle that the potential sources of knowledge and justification are ones we know to be fallible. Not only we know them to be fallible, but we know that we'll often lack the resources for determining which of their deliverances of the sources provide us with knowledge and which ones are misleading. If the normality structure is not allowed to play favourites then for large sets of individual beliefs issued by these sources our evidence will often make it (a) equally normal for any of them to be false and (b) more normal for some of them to be false than none. In those cases the smooth option will require saying that no individual belief is justified. But normic support is too fragile or too disconnected from our ordinary sense of which beliefs are justified to do justice to the intuitions behind Harman's puzzle if it tells us that acknowledging that our favourite newspaper makes mistakes that we won't spot means that we cannot acquire any justified individual belief from this paper. One could try and reject (b) by insisting that likelihood of error does not make for normality. However normic support is again too fragile or too disconnected from our ordinary sense if it requires saying that we cannot be justified in believing that we have some false perceptual belief or memory.

Can we live with that option's smooth scepticism? We think not. Granted, it allows us to believe that most of our perceptual impressions are true. As Harman remind us, we know that the sources of our testimonial knowledge make mistakes and that they'll eventually provide us with misleading evidence that is just as convincing as the genuine article. If knowing that means defeats our individual beliefs then most of our individual beliefs are defeated. But recall the problem with statistical evidence. One nice feature of the normic support view is that it purports to explain two things about evidence in the law. First, it purports to explain the non-comparative fact that the evidence we have in Prisoners is not sufficient for conviction. Second, it purports to explain the comparative fact that eyewitness testimony can be sufficient for conviction. Now consider a slightly modified version of the Eye Exam case:

Judicial Review. An informant has information that you didn't have at the time of trial and informs you that you've done very well as a judge identifying the guilty and sentencing them to jail. After years of trials, she tells you, you've made only one mistake!

Someone like Sutton might say that you shouldn't believe the one innocent defendant to be guilty and perhaps that you ought to release the one innocent prisoner while keeping the others locked up. Someone like Smith might object that this is too externalist! Fair enough. But Sutton can now reply that Smith faces a dilemma. For if he takes the spotty line, his view shares the kind of externalist feature that he objects to in Sutton's view. If he takes the smooth line, his view entails that we should not believe that any particular prisoner is guilty, only that most of them are. But that is exactly what he says we are justified in believing in Prisoners. If appropriate punishment of a prisoner requires justified belief that they are guilty, the smooth option implies that we ought to release all the prisoners. But couldn't Sutton object that this is an overreaction to the discovery that all but one of the defendants was guilty? If it is, the smooth option invalidates Smith's explanation of the comparative epistemic fact fails. For on the smooth option Prisoners and Judicial Review are epistemically alike. Hence Smith's account either faces the 'too externalist' objection or lacks a story of the contrast in Prisoners.¹⁰

There is a third option. Normality may be revised in either or both of two ways:

Normality*. A thinker has justification to believe p iff the p -relevant part of the thinker's evidence provides normic support for p .

Normality**. A thinker has justification to believe p iff the thinker's evidence provides p -relevant normic support for p .

Normality* relativizes one's evidence to each proposition; Normality** relativizes the normality ordering to each proposition.¹¹ Both allow the normic support view to endorse Justified Inconsistency while preserving its original diagnosis of Harman's paradox. On the Normality* line, we leave the doctor's report out of the evidence for each individual perceptual belief. We can then say that the evidence relevant to that belief does provide normic support for it even after the doctor's report.¹² On the Normality** line, we say that for each specific proposition, the normality ordering relevant to that proposition makes it less normal to be in error about it than about other propositions. Hence for each particular proposition, your total evidence after the doctor's report still provides relevant normic support to that proposition. Either line allows one to endorse

¹⁰ An anonymous referee raises an interesting question here: does knowability in the relevant range of cases correlate with the true value of Blackstone's ratio? Recall that Blackstone's ratio is a purported measure of the relative disvalue of jailing the innocent and letting the guilty go free, which Blackstone himself put above ten to one: "It is better that ten guilty persons escape than that one innocent suffer" (2001 [1763]). The suggestion, we take it, is that there is a true value of Blackstone's ratio and that it correlates with scrutural constraints on knowability. For instance, one might think that one known error defeats ten items of knowledge, that is, that if one knows that there are k errors in a set of n internally alike propositions then one knows at most $n-10k$ propositions in the set. Now suppose we adopt the idea defended by Moss (2018) that one should convict only the known guilty and let others go free. The structural constraint would then mean in cases like Judicial Review one should let out up to ten guilty people, but no more than ten, for each known innocent, in line with Blackstone's original ratio. While we don't immediately see why knowability would be so constrained, the proposal is worth exploring. Elsewhere we have explored another way of relating something like Blackstone's ratio with norms of judgement in these cases. Littlejohn (forthcoming) and Dutant and Fitelson (MS) defend a view on which what's rational to believe depends on the probability that the thinker's belief would amount to knowledge. The relevant probability threshold can be thought of capturing the relative disvalue of jailing somebody whose guilt is unknown versus that of letting those known to be guilty go free. The two approaches are not incompatible. One of the authors of this paper vacillates between the probable knowledge proposal and the straightforward knowledge norm. Because he still defends the view from his (2012) that innocent people are owed compensation whenever wrongfully convicted, he prefers the externalist accounts of just conviction that Moss defends, but he sees the virtue of characterising some more subjective notion to understand when an agent or a jury rationally comes to the conclusion that someone is guilty.

¹¹ Compare Holliday's (2015) "RS $\forall\exists$ " relevant alternatives models for knowledge. On these models, for each proposition there is a set of alternatives relevant to whether one knows it (rather than there being one set of alternatives that is relevant for each proposition to whether one knows it).

¹² That does not mean that Normality* rules out defeat: had the doctor said that some individual answer was false, their report would have been part of the evidence relevant to that answer.

Justified Inconsistency and hence smoothness. Both are compatible with the original normic diagnosis of Prisoners and the Lottery.

This third option accommodates a lot of constraints but at a high cost: it invalidates virtually *all* of Smith's logic of justification. Consider for instance the principle that justification distributes over conjunction: if you are justified in believing $p \& q$ you are justified in believing p and justified in believing q . This holds in Smith's models.¹³ On the Normality* line, it fails because your $p \& q$ relevant evidence need not be the same as your p -relevant evidence or your q -relevant evidence. On the Normality** line, it fails because the $p \& q$ -relevant normalcy ordering need not be the same as the p -relevant ordering or the q -relevant ordering. At most one can hope to preserve the pairwise consistency principle that one cannot be both justified in believing p and in believing $\neg p$, provided that p -relevant evidence and $\neg p$ -relevant evidence are constrained to be the same or that p -relevant normalcy and $\neg p$ -relevant normalcy are. But most of the attractive principles Smith (2016: 153) derives are otherwise lost.

3. Moorean Absurdities

One nice feature of Sutton's (2007) proposal is that it explains why it is impossible to justifiably believe Moorean absurdities like:

(1) Dogs bark but I don't know that they do.¹⁴

Sutton says that you can only justifiably believe what you know. But given that knowledge is factive and distributes over conjunction, it is impossible to know (1). For if what you know is true you cannot both know that dogs bark and know that you don't know this.

In the Lottery one's evidence does make (2) very probable:

(2) This ticket lost, but I don't know that it did.

But we know that the normic support view doesn't take this to show that the belief in (2) is justified. To the contrary, it says that a belief that one's ticket lost lacks normic support and hence is not justified. The question we raise is whether a body of evidence might provide normic support for propositions of the form $\neg p$ and I don't know p where p is *not* a lottery proposition but rather one has what looks like normic support for both conjuncts.

We know that the guiding idea here is that normic support turns on whether this belief would be true in all the most normal worlds in which the thinker's evidence holds. What we don't know is whether a proposition of the Moorean form can hold at all such worlds. So far as we can tell nothing in the system prevents that.

To see why we might expect that it's possible to have justification to believe (1), let's look at Smith's logic for justification. Note first that the following *reflexivity* schema is valid in his framework:

$$(E \& p) \Box \rightarrow p$$

If we read ' $E \Box \rightarrow p$ ' as saying that E provides normic support for p , this states that any body of evidence that includes p (which we may thus write as $E \& p$) normically supports p . For either there are $E \& p$ worlds or there are not. If not, $E \& p$ trivially supports p (Smith 2016: 137). If yes, trivially p holds at all most normal $E \& p$ worlds. Thus, $E \& p$ provides normic support for p . One consequence of the schema is:

$$(E \& p \& \sim Kp) \Box \rightarrow p \& \sim Kp$$

That is, if one's evidence includes both p and *I don't know that p* then it justifies believing their conjunction. If defenders of the normic view think that these beliefs cannot be justified, they need to say that one's evidence cannot include both p and *I don't know that p*. But we do not know what could rule out the possibility that such things could be part of a thinker's evidence. We cannot

¹³ See Smith's (2016: 144) discussion of weakening the consequent.

¹⁴ While we frame our discussion in terms of knowledge we note that we would be equally happy for it to be framed in terms of what one is in a position to know. While we think that it's not rational to believe Moorean absurdities, we acknowledge that some authors (e.g., McGlynn (2014)) think that it would be fine to believe p whilst believing that p is not something we can know.

rule this out by appeal to the (alleged) factivity of evidence. Even if only truths can be part of a thinker's evidence, this doesn't show that it's impossible for it to include both p and *I don't know that p* since these conjuncts can jointly be true. We could get the desired result if we insisted that only *known* facts can be part of a thinker's evidence. Now Smith doesn't seem particularly attracted to $E=K$ and in fact seems to hinge towards an account of knowledge in terms of evidence (Smith, 2016, chap. 7). Nonetheless, the $E \rightarrow K$ constraint would fall out of the idea that one knows p just if one's evidence safely supports p .¹⁵ But since Smith has not endorsed that account of knowledge we are not sure that this is a welcome addition.

Be that as it may, the putative connection between knowledge and evidence only blocks one route to justified Moorean absurdity. For the normic support framework validates the *rational monotonicity* rule (Smith 2016: 89, 141):

$$\begin{array}{l} E \Box \rightarrow p \\ \sim (E \Box \rightarrow \sim q) \end{array}$$

$$\hline (E \& q) \Box \rightarrow p$$

the idea is that if E provides normic support for p but not for $\sim q$, the conjunction of E and q must also provide normic support for p . Now let's ask whether evidence that provides normic support for p also provides normic support for the proposition that one knows p . Either it always does or it sometimes doesn't. Let's consider each option in turn.

On the first option, any body of evidence that provides normic support for p supports the stronger proposition that p is something the thinker knows (or is in a position to know):

$$\begin{array}{l} E \Box \rightarrow p \\ E \Box \rightarrow Kp. \end{array}$$

We think that this is too strong. Recall Williamson's (2014) unmarked clock case. Let p be the strongest proposition that a thinker knows about the position of the hand from a glance at the clock. We think that the evidence supports believing p . (It is, after all, something that the thinker knows.) We do not think, however, that the same evidence supports believing that p is something that this thinker can know. On their evidence it is very improbable that p is something they know. (Given their evidence, it is much more likely that they know something different than p .) Moreover, we do not see any reason to think that a world in which they know this particular proposition is more normal than any world in which they don't because they know something slightly different instead. It would not seem less normal for their evidence to hold and for them to know a slightly different proposition about the clock's hand.¹⁶

Let's consider the second option, the position that says that it's possible for a body of evidence to support p even if it does not support the stronger proposition that the thinker is in a position to know that p . If we assume this, we have to assume that there will be some case such that $E \Box \rightarrow p$ and $\sim (E \Box \rightarrow \sim Kp)$. Applying rational monotonicity we get:

$$(E \& \sim Kp) \Box \rightarrow p$$

That is, one's evidence supplemented with the fact that one did not know p would still provide normic support for p . That seems wrong. Recall that the normic picture of justified belief is one in which one is not justified in believing lottery propositions. On such a picture of justification, we think, we would want no want evidence that includes the fact that one does not know p to support p . Note further that by reflexivity:

$$(E \& \sim Kp) \Box \rightarrow \sim Kp$$

From that and the previous result the *agglomeration* rule (Smith 2016: 145) gives us:

¹⁵ See Smith (2016, chap. 7) on safe support. E safely supports p at world w iff all E -worlds close to w are p worlds. It follows that if E includes p then it safely supports p .

¹⁶ See Stalnaker (2015) for an opposing view. Note that even if, like Goodman (2013), one adopted a model of knowledge that blocks improbable knowing, one would still fall short of the principle that if one knows, one's evidence supports the hypothesis that one knows. On Goodman's models, knowing is compatible with it being significantly likely (though not more likely than not) that one doesn't know.

$$(E \& \sim Kp) \Box \rightarrow p \& \sim Kp$$

That is, if your evidence supports p without supporting that you know p , that that evidence supplemented with the proposition that you don't know p would justify belief in a Moorean absurdity. Again, that seems wrong. The mere fact that one's evidence supports p without supporting that one knows p shouldn't make it such that it would support a Moorean absurdity if supplemented with $\sim Kp$.¹⁷

One may suggest that if one's evidence supports p , it cannot be supplemented with $\sim Kp$, so no body evidence will in fact justify believing a Moorean absurdity. But note that this would not follow from the mere requirement that one's evidence is known. For we have merely assumed that one's evidence *supports* p , not that it *includes* it. So barring p and $\sim Kp$ from jointly entering one's evidence wouldn't suffice to preclude the existence of a case in which one's evidence includes $\sim Kp$ and yet supports p . For all the normic support view says, for instance, one may know that one does not know that a wall is red and nonetheless have evidence that provides normic support for the view that it is.

The challenge, then, is this. While we think that evidence that supports believing p need not support believing that p is known, we do think that justification for believing that p is not known should prevent p from being justifiably believed. As it stands, the normic support framework does not provide this middle ground position.

4. A Simple Fix

We have thus far examined a proposal about justification, the idea that a belief is justified when it is based on evidence that provides normic support. We raised two problems for it. We would like to offer a positive suggestion, one that would seem to solve the problems considered thus far without abandoning the idea that normality matters to justification.

We think that a more promising line to pursue would be to appeal to normality in our theory of knowledge and to appeal to knowledge in our theory of justification. We'll consider one way to appeal to normality in the theory of knowledge and three ways to appeal to knowledge in our theory of justification. First the theory of knowledge:

Normic Knowledge. One knows p iff one believes p on a basis on which one would only believe truths at cases at least as normal as one's case.¹⁸

Normic Knowledge is supposed to capture the idea that a kind of positive epistemic standing turns on whether the ways we form belief would only lead us astray in situations less normal than ours. Perception is a source of knowledge, say, because its deliverances normally lead us to form a true belief. In certain cases, however, it does not provide us with knowledge because situations in which they lead us astray are not much less abnormal than ours. As Smith observes, relative normality need not line up with likelihood: in lottery scenarios false belief is unlikely but normal. That being

¹⁷ A related puzzle is raised by Dorst (forthcoming). If one's body of knowledge includes p but not Kp , it looks like it can be augmented with $\sim Kp$ and still support p . If so, it would support conditionals of the form "If I don't know p then p ", which Dorst deems "abominable". Dorst's preferred solution is to deny the antecedent and endorse the KK principle: if one's body of knowledge includes p then it also includes Kp . As we said above, we think KK is too strong. We think the proposals sketched in the next section offer a more promising path to account for the infelicity. For instance, provided Moorean pairs of beliefs cannot be rationally believed (more on this below), the conditionals are "junk" in Sorensen's (1988) sense that one cannot come to rationally believe the antecedent without losing rational belief in the conditional. (Given the mild assumption that if it's rational to believe "If I don't know p then p " it's rational to believe "I know p " for the latter is equivalent to the material implication corresponding to the conditional, namely "I know p or p ".)

¹⁸ See Jenkins (2006), Ball (2013), Greco (2014), Stalnaker (2015), Dutant (2016) and Goodman and Salow (2018) for accounts along those lines. Views of that kind are sometimes put forward to vindicate the KK principle (Greco, Stalnaker, Goodman and Salow). This is controversial, however. In addition to Williamson's (2000) criticisms of the principle, readers should consult Goodman (2013, §3) who argues that normality conditions on knowledge are plausible independently of KK and Carter (forthcoming) who argues that the account does not support the KK principle.

said, we will be maximally hand-wavy about how to think of bases and the normality ordering. The crucial moves in this discussion don't turn out on the finer details of the account.

Now the theory of justification. We'll consider three options:

J = Potential K: One justifiedly believes p iff it is possible for an intrinsic duplicate of one to know p . (Ichikawa 2014; see also Bird 2007).¹⁹

J = Probable K: One has justification to believe p iff it is sufficiently probable on A's evidence that A is in a position to know p . (Blome-Tillmann 2015; Dutant and Fitelson forthcoming).²⁰

Ignorance is Strength: One has justification to believe p iff one is not in a position to know that one is not in a position to know p (Rosenkranz 2018; see also Lenzen 1978, Stalnaker 2006, Dutant forthcoming).²¹

On these accounts justification does not require knowledge. Combined with Normic Knowledge, they say that justification does not require normic support but rather that one's belief be suitably related to a belief that receives normic support. According to the first proposal, a belief is justifiably held even if it does not constitute knowledge, provided that a counterpart holds a belief that does constitute knowledge. According to the second, the crucial question is whether the probability on a believer's evidence that they know is sufficiently high. And according to the third, a belief is justified when the believer cannot know that they do not know. (If they know, they cannot know that they do not know, so knowledge entails justification. Justification does not entail knowledge, only that the thinker is not in a position to know that they do not know some target proposition.)

Each of these accounts has resources to explain Harman's contrast. In a typical lottery case, one is not internally like someone who knows, it is arguably not probable on one's evidence that one knows, and one is arguably in a position to know that they do not know. In a typical testimony case, however, one is arguably internally like someone who knows, it is probable on one's evidence that one knows, and one is not in a position to know that they do not know. The accounts also have the resources to vindicate Justified Inconsistency. For we may think that the eye doctor's report is not enough to prevent you from knowing when you did perceive things well. If so, for each of your perceptual beliefs you are like someone who knows, it is probable on your evidence that you know, and you are not in a position to know that you do not know. Of course

¹⁹ Bird (2007) doesn't frame the view in terms of internal duplicates but rather in terms of subjects sharing their mental (incl. factive) states before forming the target belief. Littlejohn (2018) discusses some difficulties preface cases present for these views when they're combined with E=K and related accounts of evidence (e.g., that evidence consists of basic knowledge). Conditionalising on such evidence can make it hard to see how each belief in the relevant inconsistent sets can be justified. This problem is similar in some ways to the problem that Praolini (forthcoming) raises for epistemologists who accept certain closure principles.

²⁰ Note that while this view handles Eye Exam and Judicial Review reasonably well, it will have to deny that knowledge is sufficient for justification if there are cases of improbable knowing (Williamson 2014) or cases in which knowing is compatible with an insufficiently high probability that one knows (as e.g. in Goodman 2013). For a defence of 'unreasonable knowledge', see Lasonen-Aarnio (2010). In Littlejohn and Dutant (forthcoming), we defend a view of rational belief and the defeat of rationality on which *ex ante* rationality is characterised in terms of the probability of being in a position to know. We think that this approach provides a much more straightforward account of higher-order defeat and negative self-appraisal defeat (i.e., the defeat associated with judgments about what can be known, what can be rationally believed, what the evidence supports) than more familiar truth-centric approaches to rationality.

²¹ Lenzen (1978) and Stalnaker (2006) use the idea as an account of *belief* rather than justified belief (though their notion of belief is idealized). Dutant (forthcoming) puts forward a related account on which justification to believe p doesn't merely require that p is known at *some* epistemically possible case but that p is known at *all* epistemically possible cases that are "best" along some dimension. On all these proposals Moore-paradoxical pairs of beliefs cannot be justified. On the Rosenkranz, Lenzen and Stalnaker proposals the result requires the luminosity of justification, for substituting $\sim K \sim K$ for J in $Jp \rightarrow \sim J \sim Kp$ gives us $\sim K \sim Kp \rightarrow \sim \sim K \sim K \sim Kp$. On Dutant's proposal that follows from factivity alone, since no matter what the best epistemically possible cases are they are not cases where one knows p and that knows that one doesn't know p .

you are not internally like one who knows *all* of your perceptual answers *and* that one of them is incorrect, it is not probable that you know *all* the answers and you are in position to know that you do not know *all* of them. But provided we apply the accounts to each belief separately that does not prevent individual beliefs from being justified. Ditto for Judicial Review. Your justified beliefs are neither consistent nor closed, but that may be defensible. Indeed, we think that is the correct attitude in such a case. Finally, each of the accounts has the resources to explain Moore's paradox. No one can know both p and that they do not know p , so one cannot be internally like someone who does, is it not probable that one knows the conjunction, and one is always in a position to know that one does not know the conjunction.

Smith's normic support theory of justification, as well as Sutton's J=K proposal, force us to say that an agent's justified beliefs have to be consistent. By our lights, this means that these views confront a dilemma. They might be too externalist (i.e., they predict differences in justification on grounds that are inaccessible to the thinker) or too sceptical (i.e., they deny us too much justification in cases like Eye Exam and Judicial Review). When we combine Normic Knowledge with one of our knowledge-centred theories of justification, we avoid this dilemma. In Eye Exam and Judicial Review, each of the particular beliefs that constitute the inconsistent set are promising candidates for knowledge. Although it is certain that not all these beliefs constitute knowledge, each of them meets the right-hand side of the conditions spelled out in our proposed theories of justification. We see this as a major advantage of these approaches.²²

References

- Adler, Jonathan. 2002. *Belief's Own Ethics*. MIT University Press.
- Ball, Brian. 2013. Knowledge is Normal Belief. *Analysis* 73: 69-76.
- Bird, Alexander. 2007. Justified Judging. *Philosophy and Phenomenological Research* 74: 81-110.
- Blackstone, William. 2001 [1763]. *Blackstone's Commentaries on the Laws of England*. Wayne Morrison (ed.) Cavendish Publishing.
- Blome-Tillmann, Michael. 2015. Sensitivity, Causality, and Statistical Evidence in Courts of Law. *Thought: a Journal of Philosophy* 4: 102-12.
- Buchak, Laura. 2013. Belief, Credence, and Norms. *Philosophical Studies* 169: 1-27.
- Carter, Sam. Forthcoming. Higher Order Ignorance Inside the Margins. *Philosophical Studies*.
- Christensen, David. 2004. *Putting Logic in its Place*. Oxford University Press.
- Cohen, Stewart. 1984. Justification and Truth. *Philosophical Studies* 46: 279-95.
- Dorst, Kevin. Forthcoming. Abominable KK failures. *Mind*.
- Dutant, Julien. 2016. How to be an Infallibilist. *Philosophical Issues* 26: 148-171.
- Dutant, Julien. Forthcoming. Knowledge-First Evidentialism about Rationality. In F. Dorsch & J. Dutant (ed.) *The New Evil Demon Problem*. Oxford University Press.
- Dutant, Julien and Fitelson, Branden. MS. Knowledge-Centered Epistemic Utility Theory.
- Dutant, Julien and Littlejohn, Clayton. Forthcoming. Defeaters as Indicators of Ignorance. In J. Brown and M. Simion (ed.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Easwaran, Kenny and Fitelson, Branden. 2015. Accuracy, Coherence, and Evidence. In T. Gendler and J. Hawthorne (ed.), *Oxford Studies in Epistemology*, Vol. 5. Oxford University Press: 61-96.
- Evnine, Simon J. 1999. Believing Conjunctions. *Synthese* 118: 201-27.
- Foley, Richard. 1992. The Epistemology of Belief and the Epistemology of Degrees of Belief. *American Philosophical Quarterly* 29: 111-121.
- Gardiner, Georgi. Forthcoming. Evidentialism and Moral Encroachment. In K. McCain (ed.), *Believing in Accordance with the Evidence: New Essays on Evidentialism*. Springer.

²² The authors would like to thank Alexander Bird, Branden Fitelson, Dan Greco, John Hawthorne, Francesco Praolini, Sven Rosenkranz, Sherri Roush, Bernhard Salow, Martin Smith, and Tim Williamson for discussing these issues with us. We would also like to thank an anonymous referee for this journal for their helpful comments.

- Gerken, Mikkel. Forthcoming. Truth-Sensitivity and Folk Epistemology. *Philosophy and Phenomenological Research*.
- Goodman, Jeremy. 2013. Inexact Knowledge without Improbable Knowing. *Inquiry: An Interdisciplinary Journal of Philosophy* 56: 30-53.
- Goodman, Jeremy and Salow, Bernhard. 2018. Taking a Chance on KK. *Philosophical Studies* 175: 183-96.
- Greco, Dan. 2014. Could KK be OK? *Journal of Philosophy* 111: 169-97.
- Holliday, Wesley H. 2015. Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. *Journal of Philosophical Logic* 44: 1–62.
- Huemer, Michael. 2006. Phenomenal Conservatism and the Internalist Intuition. *American Philosophical Quarterly* 43: 147-158.
- Ichikawa, Jonathan Jenkins. 2014. Justification is Potential Knowledge. *Canadian Journal of Philosophy* 44: 184-206.
- Jenkins, Carrie. 2006. Knowledge and Explanation. *Canadian Journal of Philosophy* 36: 137-164.
- Lasonen-Aarnio, Maria. 2010. Unreasonable Knowledge. *Philosophical Perspectives* 24: 1-21.
- Leitgeb, Hannes. 2017. *The Stability of Belief*. Oxford University Press.
- Lenzen, Wolfgang. 1978. Recent Work on Epistemic Logic. *Acta Philosophica Fennica* 30: 1–220.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, Clayton. 2018. Objectivism and Subjectivism in Epistemology. In V. Mitova (ed.), *The Factive Turn*. Cambridge University Press, 142-60.
- Littlejohn, Clayton. Forthcoming. Truth, Knowledge, and the Standard of Proof in Criminal Law. *Synthese*.
- McDowell, John. 1998. *Meaning, Knowledge, and Reality*. Harvard University Press.
- McGlynn, Aidan. 2013. Believing Things Unknown. *Notis* 47: 385–407.
- Makinson, D.C. 1965. The Paradox of the Preface. *Analysis* 25: 205-7.
- Moss, Sarah. 2018. *Probabilistic Knowledge*. Oxford University Press.
- Nelkin, Dana. 2000. The Lottery Paradox, Knowledge and Rationality. *Philosophical Review* 109: 373-409.
- Pollock, John. 1974. *Knowledge and Justification*. Cornell University Press.
- Pollock, John. 1986. The Paradox of the Preface. *Philosophy of Science* 53: 246-258.
- Praolini, Francesco. Forthcoming. No Justificatory Closure Without Truth. *Australasian Journal of Philosophy*.
- Pryor, Jim. 2000. The Skeptic and the Dogmatist. *Notis* 34: 517-49.
- Redmayne, Michael. 2008. Exploring the Proof Paradoxes. *Legal Theory* 14: 281–309.
- Rosenkranz, Sven. 2018. The Structure of Justification. *Mind* 127: 309-338.
- Roush, Sherrilyn. 2010. Optimism about the Pessimistic Induction. In P.D. Magnus and J. Busch (ed.), *New Waves in Philosophy of Science*. Palgrave-MacMillan, pp. 29-58.
- Ryan, Sharon. 1991. The Preface Paradox. *Philosophical Studies* 64: 293-307.
- Smith, Martin. 2016. *Between Probability and Certainty*. Oxford University Press.
- Smith, Martin. 2018. When Does Evidence Suffice for Conviction? *Mind* 127: 1193-1218.
- Sorensen, Roy. 1988. Dogmatism, Junk Knowledge, and Conditionals. *The Philosophical Quarterly* 38: 433-454.
- Stalnaker, Robert. 2006. On the Logics of Knowledge and Belief. *Philosophical Studies* 128: 169-199.
- Stalnaker, Robert. 2015. Luminosity and the KK Thesis. In S. Goldberg (ed.) *Externalism, Self-Knowledge and Skepticism*. Cambridge University Press.
- Sutton, Jonathan. 2007. *Without Justification*. MIT University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- Williamson, Timothy. 2013. Gettier Cases in Epistemic Logic. *Inquiry* 56: 1–14.
- Williamson, Timothy. 2014. Very Improbable Knowing. *Erkenntnis* 79: 971-999.
- Worsnip, Alex. 2016. Belief, Credence, and the Preface Paradox. *Australasian Journal of Philosophy* 94: 549-62.