Normic Support: A Risky Variant

Andrew Peet

May 16, 2022

1 Introduction

You possess a ticket for a large lottery. The draw has taken place but you have not yet seen the results. You know that there is guaranteed to be a winner, and that the chances of winning are minuscule. Are you justified in believing that your ticket lost? Many say no. There are a number of reasons for this.

Firstly, it seems clear that you cannot know that you have lost. And if you know that you don't know that p, surely you cannot rationally believe that p. As Littlejohn (2012) points out, if you were justified in believing that you will lose the lottery, then it must be rational for you to believe 'I will lose the lottery and I don't know that I will lose the lottery'. But this is Moore paradoxical. Secondly, as Nelken (2000) suggests, it is seems irrational to believe that p when you know there is no explanatory connection between the fact that p and your belief that p. If you believe that your ticket is a loser on the basis of the probabilities alone then your belief will bear no explanatory connection to the ticket's actually being a loser. So, your belief will be irrational.

Finally, the main reason to hold that you *are* justified in believing your ticket will lose is the assumption that one's level of justification for believing a proposition p is identical to the probability of p on one's evidence. On this assumption, the high probability that your ticket is a loser grants a high degree of justification to your belief that it is a loser. However, there is reason to be suspicious of this assumption. Take the following example from Smith (2010):

> **Computer Screen** Suppose that I have set up my computer such that, whenever I turn it on, the colour of the background is determined by a random number generator. For one value out of one million possible values the background will be red. For the remaining 999 999 values, the background will be blue. One day I turn on my computer and then go into the next room to attend to something else.
>
> In the meantime Bruce, who knows nothing about how my computer's background colour is determined, wanders into the computer room and sees that the computer is displaying a blue background. He comes to believe that it is. Let's suppose, for the time being, that my relevant evidence consists of the proposition that (E_1) it is 99.9999% likely that the

computer is displaying a blue background, while Bruce's relevant evidence consists of the proposition that ((E_2) the computer visually appears to him to be displaying a blue background. Smith, 2010, 13-14.

Intuitively, Bruce's justification for believing that the background is blue is greater than Smith's. Smith's attitude would be best described as a presumption. He should not outright assert that the background is blue, he should report only that it is highly likely that the background is blue. Bruce, on the other hand, can outright assert that the screen is blue. His belief is no mere presumption. Yet, although it is unlikely, Bruce could still be mistaken. For example, he could be hallucinating. The probability of Bruce being mistaken is extremely low. Yet it is surely lower than the probability of Smith being mistaken. After all, it is 99.9999% likely that Smith's belief is correct. So, if our degree of justification for a proposition p was simply the probability of p on our evidence we would expect Smith's belief to be more justified than Bruce's.

A number of approaches to justification entail that we are not justified in believing we will lose the lottery. Here my focus will be on one: Smith's (2010, 2016) normic support view.[1] Smith tells us that a body of evidence E supports p iff p is true at all the most normal worlds consistent with E. Assuming that one's belief that p is justified only if one's evidence supports p, this entails that a justified belief must be true at all the most normal worlds consistent with the believer's evidence. Normality here is thought of in non-statistical terms. More specifically, abnormal events are those that cry out for explanation. Normal events do not. The more normal a world is the more idealized it is (Smith (2007)). If, given one's evidence, it would not cry out for explanation for one's belief to be false, then one's belief is not justified. One's belief that the lottery ticket is a loser is like this: given one's evidence, the ticket's being a winner would not cry out for explanation. It could just so happen that this is the case. So, the ticket is a winner at one of the most normal worlds consistent with one's evidence, and you are not justified in believing that it is a loser.

Normic support has a number of problems. In particular, it entails that we can be justified in believing arbitrarily improbable propositions, it fails to explain the relationship between justified belief and justified credence, and it allows the possibility Moore paradoxical justified beliefs such as 'p and I don't know that p' (Littlejohn & Dutant (2020)).[2] Two of these problems mirror

[1] Knowledge first views of justification constitute a prominent alternative to the normic support approach (see Sutton (2007), Ichikawa (2014), Blome-Tillman (2015), and Rosenkranz (2018), Littlejohn & Dutant (2020)). The view I go on to advocate is close in important respects to the knowledge first approach. In particular, it implies that, on certain reasonable assumptions about knowledge, a subject is justified in believing that p if and only if they would be in a position to know that p in normal circumstances. However, it does not analyze justification in terms of knowledge.

[2] These are not the only challenges to have been raised for normic support. They are merely the challenges I take to be most pressing. See Backes (2019a) for a further objection, and Smith (Forthcoming a) for a response. There are a number of objections that turn on normic support's validation of multi premise closure (for example Backes (2019b), and Praolini (2019). I will not be concerned with these issues here as the version of normic support I advocate invalidates multi premise closure.

a problem in the theory of risk that can be resolved by embracing the hybrid modal/probabilistic account of risk developed by Peet (XXXX). We can combine this theory of risk with the normic support picture of justification as follows: a body of evidence E supports p if and only if there is no risk of p being false at the most normal worlds consistent with E. The resultant approach entails that lottery beliefs are not justified, and that we cannot be justified in believing highly improbable propositions. It also entails that one's evidence cannot contain both p and the proposition that one does not know that p. And it blocks the derivation of Moore paradoxical sentences. Unlike knowledge first approaches (the main rival to normic support) it does not analyze justification in terms of knowledge. However, it does imply, at least on certain natural assumptions about knowledge, that one will be justified in believing p iff one is in a position to know that p at the most normal worlds consistent with one's evidence. I will start by explaining why the standard normic support approach entails that we can be justified in believing arbitrarily improbable propositions.

2 Prefaces and Probability

A famous historian writes an intimidatingly long, and incredibly well researched book. The historian believes every claim in the book, and they are justified in doing so. After all, meticulous research has gone into every single claim. However, they are also aware that similar books, even if well researched, have always contained errors. They realize that the overall probability that every claim in the book is true is vanishingly small. So, they write in the preference 'no doubt there are some false claims in this book'. This seems paradoxical. On the one hand they believe every claim in the book, and they are justified in doing so. On the other hand, it appears that they also believe that the conjunction of these claims is false. And this is contradictory. After all, the only way for the conjunction to be false would be for one of the claims in the book to be false (Makinson (1965)).

Perhaps the most common response to the preface paradox is to simply accept the possibility of justified inconsistent belief sets. This is fairly natural if we adopt a probabilistic threshold view of justification (Foley (1992)): suppose that one is justified in believing that p whenever the probability of p conditional on one's evidence surpasses some threshold (say, 0.7). Suppose that all the claims in the book are independent, and the probability of each claim on the historian's evidence is 0.9. Then the historian will be justified in believing every claim. However, if there are a great many independent claims in the book, then the probability of them all being true will be very low. The probability that their conjunction is false will surpass our threshold for justification (0.7). So, our historian will also be justified in believing that the conjunction is false.[3]

[3] Littlejohn and Dutant (2020) argue that knowledge-first approaches (other than Sutton (2007)) also allow for justified inconsistent belief sets. Suppose that, before writing the preface, an all knowing oracle tells our historian that all but one of the claims in the book are true. In this case, our historian knows that there is one false claim in the book. So, they know (and

Many embrace the conclusion that our historian's inconsistent belief set is justified. However, it is not clear that they ought to. The basis for our historian's belief in the falsity of the conjunction is simply the fact that the conjunction has an extremely low probability. However, if we accept that lottery beliefs are not justified it is not clear why we should accept that our historian's belief is justified. After all, high probability is not sufficient for justification. And it certainly seems prima facie unattractive to admit the possibility of justified inconsistent belief sets. As Adler (2002, p 200) puts it: 'it seems incoherent to hold, as it would be incoherent to assert, that each of a set of one's beliefs is true including the belief that at least one of them is not.'[4] This is not to say that our preface writer should arrogantly deny their own fallibility. They can easily express their fallibility by stating 'there is a high probability that this book contains at least one false claim'. On can consistently acknowledge this without outright believing that the the book contains at least one false claim. After all, one can coherently acknowledge that a proposition has a high probability without outright believing that claim (as we see from lottery cases).

Still, admitting justified inconsistent belief sets does seem to be better than the alternative: maintaining that our historian is justified in believing the conjunction of the claims in their book. This would force us to accept that one can justifiably believe propositions that are arbitrarily improbable conditional on one's evidence. After all, we can construct preface cases in which our historian's book contains arbitrarily many equally well supported claims. This entails that one could be justified in believing p despite knowing that the probability of p conditional on one's evidence is, say, 0.00000000000001. This is an unacceptable result. Yet, it is exactly the result we get on the normic support

are therefore justified in believing) that the conjunction of the claims in their book is false. Littlejohn and Dutant suggest that this is consistent with their still knowing the vast majority of the claims in their book: for each claim, our historian will have a possible internal duplicate who knows that claim. Moreover, for each claim, the probability that our historian knows that claim is high. And, our historian is not in a position to know of any claim in the book that they do not know that claim to be true. So, on most knowledge first approaches to justification, they will be justified in believing every claim in the book (Ichikawa (2014), Blome-Tillman (2015), and Rosenkranz (2008)). However, it is questionable whether our historian really does retain their knowledge of any of the individual claims in the book. As Smith (Forthcoming a) points out, the position of our historian is very similar to that of our lottery ticket holder. Our lottery ticket holder knows of each ticket that there is a high probability of its being a loser. However, they also know that one ticket is definitely a winner. Likewise, our historian knows that the chances of each claim being false are very low. However, they also know with certainty that one such claim is false. So it is natural to conclude that, like our lottery ticket holder, they no longer know any individual claim in the book. Rather, they merely know that each claim is highly probable. If this is right (as I suspect it is) then the knowledge-first approach will not imply the possibility of justified inconsistent belief sets. This also marks an important dis-analogy with standard preface cases: in standard preface cases the author does not *know* that there is a false claim in their book. They merely know that there is a high probability of at least one claim being false. It is, thus, more plausible to maintain that the standard preface writer remains justified in believing the individual claims in their book than our oracle informed preface writer.

[4]For further arguments against justified inconsistent belief sets see Ryan (1991), Evnine (1999), Adler (2002), and Leitgeb (2017).

approach to justification.[5]

The normic approach tells us that a belief that p is justified iff it is true at all the most normal worlds consistent with the believer's evidence. Suppose our historian's book only contains three claims: p, q, and r. The historian is justified in believing p, so p is true at all the most normal worlds consistent with our historian's evidence. Likewise for q, and likewise for r. But this means that p, q, and r are all true at all the most normal worlds consistent with our historian's evidence. And this entails that the proposition $\ulcorner p\&q\&r \urcorner$ is true at all of the most normal worlds consistent with our historian's evidence. So, our historian's evidence will support $\ulcorner p\&q\&r \urcorner$. Moreover, the only way for $\ulcorner \neg(p\&q\&r) \urcorner$ to be true is for one of p, q, or r to be false. So there are no worlds consistent with our historian's evidence at which $\ulcorner \neg(p\&q\&r) \urcorner$ is true. Thus, they cannot be justified in believing it. This generalizes. No matter how many claims are made in our historian's book, and no matter how improbable their conjunction, as long as each claim is justified their conjunction will also be justified. So, on the normic support approach to justification, we can be justified in believing arbitrarily improbable propositions.

This is absurd, and I believe it is one aspect of a more general problem facing the normic support approach: Lottery cases show us that we cannot equate justification with probability - a high probability of p conditional on one's evidence is not sufficient for one to be justified in believing p. However, there are still clear and important connections between justification and probability. For example, *typically* when p has a high probability conditional on one's evidence one will be justified in believing that p. This is no coincidence. Yet, the normic support approach completely divorces justification from probability. In doing so it leaves us unable to explain why we are typically justified in believing that p when p has a high probability on our evidence.

So, the probabilistic threshold account of justification seemingly forces us to accept the possibility of inconsistent belief sets. The normic support approach on the other hand forces us to accept that we can be justified in believing arbitrarily improbable propositions. Neither possibility is attractive. It would be better if we could hold that our preface writer is justified in believing every claim in their book, and that they are neither justified in believing the conjunction of these claims, nor in believing the conjunction to be false. To achieve this result I will modify the normic support approach as follows: a body of evidence E supports p iff there is no risk of p being false at any of the most normal worlds consistent with E. In order to understand how this works we need to investigate

[5] Smith (Forthcoming b) argues that it is hard to avoid this result as, at least in some cases, believing p and q individually one automatically brings it about that one believes $\ulcorner p\&q \urcorner$. He challenges the opponent of improbable justified belief to show that there are no such cases. This is not the place to provide a full view of belief. But I suspect that Smith's challenge can be met by a number of approaches to belief. I am particularly sympathetic to the view that to believe p is to be disposed to treat p as a premise in practical reasoning (and perhaps to be disposed to mentally affirm p). On such a view believing p and q individually does not automatically bring it about that one believes $\ulcorner p\&q \urcorner$. After all, one can be disposed to treat both p and q as premises in practical reason, and to mentally assent to them, without being disposed to do so with respect to $\ulcorner p\&q \urcorner$.

the concept of risk.

3 Probability & Risk

Risk is typically thought of in probabilistic terms. It is natural to assume that there is a (relevant) risk of p whenever the probability of p surpasses some threshold. This natural approach to risk is obviously very close to the probabilistic conception of justification. Indeed, we can think of probabilistic conceptions of justification as holding that one's belief that p is justified when the (probabilistically construed) risk of one's being mistaken is sufficiently low. Smith (2010, 2016) even labels this the 'risk minimization' conception of justification.

However, recent approaches to risk (Williamson (2009), Pritchard (2015), and Ebert, Smith and Durbach (2020)) have questioned the centrality of probability to the concept of risk. They identify a number of problems for the probabilistic approach.

The first problem is that probabilistic approaches fail to validate certain intuitively valid patterns of reasoning. Consider the following argument:

1. There is no (relevant) risk of p.

2. There is no (relevant) risk of q.

$\therefore$ There is no (relevant) risk of $\ulcorner p \vee q \urcorner$

This reasoning seems obviously valid. Moreover, as Ebert, Smith, and Durbach point out, it is exemplified by our practices of *de minimis* risk assessment, and our norms of legal fact finding. However, it is invalid on probabilistic conceptions to risk. Suppose that the probabilistic threshold for relevant risk is n. Furthermore, suppose that the probability of p is just marginally below n, and the probability of q is just marginally below n. Finally, suppose that p and q are independent. If this is the case then there will be no relevant risk of p, and there will be no relevant risk of q. However, there may well be a relevant risk of $\ulcorner p \vee q \urcorner$. After all, the probability of $\ulcorner p \vee q \urcorner$ could easily surpass n.

The second problem is that there can be a relevant risk of p even when p has an extremely low probability. Suppose a nuclear bomb is set to go off if a particular lottery ticket wins. In this situation there is a relevant risk of the bomb going off. Yet, the probability of its going off is extremely low.[6]

So, the probabilistic threshold account of risk faces major problems. Moreover, these problems are very similar to those faced by the probabilistic threshold account of justification. Two alternatives to the probabilistic account of risk have been presented. Both have structural similarities to the normic support conception of justification. Modal approaches to risk (Williamson (2009), and Pritchard (2015)) tell us that there is a relevant risk of p whenever there is a

[6]For further objections to the probabilistic account see Williamson (2009) and Pritchard (2015). The view of risk developed by Peet (XXXX) and employed here captures Williamson & Pritchard's data, but I will leave it to one side here for reasons of space.

sufficiently nearby world at which p is true. There is no relevant risk of p if there is no sufficiently nearby world at which p is true. Normic approaches to risk (Ebert, Smith, and Durbach (2020)) tell us that there is a relevant risk of p if and only if there is a sufficiently normal world at which p is true. We can think of the normic support picture of justification as telling us that a belief that p is justified iff, given the evidence, there is no (normic) risk of p being false.

Unfortunately, the modal and normic accounts of risk face a problem, and that problem is very similar to that faced by normic approaches to justification: they allow that there can be an arbitrarily high probability of p without there being a relevant risk of p. Consider our famous historian again. Suppose that their book consists in a long series of predictions about the future. As a result, the truth value of each claim in the book is metaphysically unsettled. Nonetheless, our historian has very strong evidence for each claim. So, we might suppose, there are no normal worlds at which any of the claims are false (we can also suppose that there are no nearby worlds at which they are false). However, the probability of the conjunction of the claims in the book will be extremely small. Suppose that there is a 0.0001 chance of every prediction being accurate. Furthermore, suppose that a huge bomb will go off if any of the predictions turn out to be false. The probability of the bomb going off is extremely high. It seems clear that there is a serious risk of the bomb going off. But the normic and modal accounts are unable to accommodate this.[7]

This is one aspect of a more general problem for modal and normic approaches to risk: they completely divorce risk from probability. Yet, despite plausible arguments to the effect that risk is not *just* a matter of probability, the two are still clearly related. And this is something we need to account for.

4 The Hybrid View of Risk

Peet (XXXX) argues that these problems can be solved by adopting an account of risk that combines the modal and probabilistic approaches.[8] Peet's account makes a number of simplifying assumptions. Firstly, we assume that the set of nomologically possible worlds is finite. Secondly, we assume that the physical chance of p is proportional to the number of worlds at which p is true. So, if $P(p) = n$ then p is true at $(n \cdot 100)\%$ of worlds. With these assumptions in place, the basic idea is as follows: there is a risk of p if and only if there is a sufficiently nearby world at which p is true. The level of risk (which we can represent on a [0,1] scale) is is a function of the proportion of sufficiently nearby worlds at which p is true. For example, if 50% of the sufficiently nearby worlds are p worlds, then p's level of relevant risk will be 0.5.

[7] Ebert, Smith, and Durbach (2020) adopt a pluralist account of risk according to which there are both normic and probabilistic conceptions of risk. They would hold that in this scenario there is a probabilistic risk, but not a normic risk. In principle a similar approach could be taken to epistemic justification - we could hold there are both probabilistic and normic conceptions of justification. I will pursue a monist approach here.

[8] My exposition in this section draws heavily on Peet (XXXX)'

So far this is similar to the standard modal account of risk. To capture the relationship between probability and risk, we need to build a little more into this framework. Firstly, we need say what counts as a *sufficiently* nearby world. And in order to do this we need to draw a distinction between what Peet calls 'primary risk events' and 'sub-risk events'. The basic idea here is that when we engage in risk assessment there will usually be some primary risk we are concerned with (say, p). We will determine whether there is a relevant risk of p by investigating the ways in which p could occur. The ways in which p could occur are p's 'sub-risk events'. With this distinction in hand, Peet's proposal is that the higher the probability of the primary risk event, the more p possibilities count as 'sufficiently close'. The idea being that, as Peet puts it, 'if the probability of a risk event is high, then otherwise distant possibilities become more pressing and worthy of consideration: they become more relevant' (Peet (XXXX)).

This can be modeled by assuming that there is a fixed set of worlds - what Peet calls the 'base sphere' that are always relevant when considering risk. If the probability of the primary risk event is high enough then we add p possibilities from more distant spheres to our base sphere. The more probable p is, the more distant the p possibilities we add. Peet calls the resultant set of worlds the set 'relevant worlds'. There is a risk of p if there is a relevant world at which p is true. And p's degree of risk is a function of the proportion of relevant worlds at which p is true.

This set of worlds will remain fixed whether we are directly assessing the risk of our primary risk event p, or its sub risk events p_1, p_2, and p_3. That is, when the primary risk event we are concerned with is p, and we are trying to determine whether there is a relevant risk of p occurring in the manner p_1, the set of worlds we will look at is determined by the probability of p. If p_1 occurs within this set of worlds there will be a risk of p_1, and therefore a risk of p.

This gives us what we need to vindicate checklist reasoning: Suppose p is our primary risk event, and p_1, p_2, and p_3 are p's sub-risk events. The probability of p will determine the sphere of worlds relevant to the assessment of p_1, p_2, and p_3. Suppose we determine that there are no p_1 worlds within this sphere (and so no relevant risk of p_1), no p_2 worlds within this sphere (and so no relevant risk of p_2), and no p_3 worlds in this sphere either (and so no relevant risk of p_3). Then there will be no p worlds within the sphere. So, there will be no relevant risk of p.[9]

It also allows that there can be a risk of p even when the probability of p is extremely low: when p has a low probability the set of relevant worlds will be small, perhaps not even surpassing the base sphere. However, if p is nonetheless true in one of these worlds there will still be a risk of p. This is what we find in the lottery bomb case discussed above.

However, to fully capture the relationship between risk and probability we need to build in some assumptions regarding the distribution of worlds through-

[9]The situation becomes more complex when there are multiple primary risk events at issue in a context. See Peet (XXXX) for discussion.

out modal space.[10] We will assume that distance is a measure of similarity, and that the similarity of a world to actuality is a function of the number of ways in which it differs from the actual world. To model this, Peet suggests that we think of worlds as series of 1s and 0s of equal length. Where n is the number of spheres, the number of possible worlds will be 2^{n}. And, supposing that we have S spheres of possibilities, we can determine the number of worlds within the nth sphere as $\sum_{k=0}^{n} \binom{S}{k}$. We will assign each sphere a number between 1 and 0 such that the number of the nth sphere is determined by dividing 1 by the total number of spheres, and multiplying the result by n. We will describe this number as the sphere's 'distance'. The actual world, being maximally similar to itself, will have a distance of 0. The worlds most dissimilar to actuality will have a distance of 1.

This gives us what we need to give an account of how the probability of a primary risk event determines which worlds count as sufficiently nearby. Peet outlines two pictures of this relationship, one simple, and another slightly more complex. The more complex picture allows us to model the stakes sensitivity of risk attributions, and allows for more plausible results, especially when different assumptions are made about the structure of modal space. However, for our purposes here the simple view will suffice. According to this picture, we add to the base sphere any p possibilities from the sphere of worlds that most closely approximates the probability of the primary risk event.

With this picture in hand, Peet gives us a procedure for calculating the minimal primary risk event probability required for there to be a guaranteed risk of p:

> "If we know that E has a probability of n then we know that $n\%$ of the total worlds are E worlds. For a given number of worlds this will allow us to calculate the total number of E worlds in modal space. From its distance we will be able to determine the total number of worlds within our base sphere. And for a given E probability we'll be able to calculate the total number of worlds within the E-probability sphere (which will always include all worlds within the base sphere). From here we can calculate the minimum possible number of E worlds within the set of relevant worlds. First, we assume that all (or as many as possible) of worlds not within the E-probability sphere are E worlds. This will allow us to calculate the minimum possible number of E worlds within the E-probability sphere. From here we can calculate the maximum possible number of $\neg E$ worlds within the base sphere, and add to this the minimum number of E worlds within our E-probability sphere. This will give us the total number of worlds within the set of relevant worlds, and will allow us to calculate the minimum possible proportion of E

[10] These precise assumptions are not essential. As Peet (XXXX) explains, the view is workable on various plausible conceptions of the distribution of worlds throughout modal space (although it yields implausible results if, for example, almost all worlds are maximally dissimilar to the actual world). However, these assumptions are useful for presentational purposes.

worlds within this set to give E's minimum guaranteed degree of risk. We can go through the same procedure when E is a sub-risk event. However, the E-probability sphere will not be determined by the probability of E. Rather, it will be determined by E's primary risk event." Peet (XXXX), pp xxxx

For our purposes the most important implications of this picture are that as long as the probability of p is higher than 0.5 there is guaranteed to be a risk of p. And, for all values above 0.5, the higher the probability of p the higher the guaranteed risk of p. If p has a very high probability then there is guaranteed to be a high risk of p. This allows us to ensure that there is a high degree of risk in the preface bomb case.

5 Risk and Justification

So, similar problems arise when we divorce risk and justification from probability. We have just seen how Peet's (XXXX) hybrid modal/probabilistic approach helps with risk. Here I will show that it also helps with justification. The basic idea is that a subject is justified iff, given their evidence, there would not normally be any risk of error. This is an intuitively appealing hypothesis. After all, suppose that all the *most* normal worlds consistent with E are p worlds, but there is, at all the most normal worlds consistent with E, a very high risk of p being false. Smith's account tells us that E supports p in such circumstances. But this seems wrong. We are surely not justified if we form our beliefs in such a way that there would normally be a high risk of error. So, if there would normally be a (relevant) risk of error given one's evidence, one's belief is not justified.

More precisely, the picture is as follows: a body of evidence E supports a proposition p if there is no relevant risk of $\neg p$ (considered by itself) at the most normal worlds consistent with E.[11] So, p must be true not only at all the most normal worlds consistent with E, but also a sufficiently wide sphere of worlds around the most normal worlds consistent with E (these worlds must also be consistent with E). How wide of a sphere we consider will depend on the probability of p in the manner outlined above.[12] If p has a high probability conditional on E, then the probability of $\neg p$ conditional on E will be low, meaning we only need to investigate a narrow sphere of worlds around the most normal worlds consistent with E. However, even when p has a high probability conditional on E, E won't always support p. For example, if $\neg p$ would only require a minimal deviation from normality (or no deviation from normality at all), despite its low probability, then E won't support p. This is exactly what we find in lottery cases.

[11] Importantly, when assessing whether a belief that p is justified, the risk of p being false will be the primary risk at issue.

[12] I assume that $P(p|E)$ is the same at the actual world as it is at the most normal worlds consistent with E. If these come apart then we will say that the risk of p at w is dependent on the probability of p at w.

On the other hand, if p has a low probability conditional on E, then the probability of $\neg p$ will be high, meaning that we will have to consider a wide sphere of worlds. As outlined above, there is guaranteed to be some relevant risk of $\neg p$ when p has a low probability. So, if p has a low probability conditional on E then E will not support p. We will cannot be epistemically justified in believing improbable propositions.

Consider the preface paradox again: Each claim in our historian's book has a high probability of being true, and each claim is true at the most normal worlds consistent with the evidence. Since the probability of each claim is high, each claim will only have to be true throughout a narrow sphere of normal worlds in order for the claim to be justified. So, assuming each claim is well researched, each claim will be justified. However, the probability of the conjunction of the book's claims is vanishingly small. So, in order for it to be justified it will have to be true throughout a very large sphere of worlds, including worlds that are highly abnormal. Given the high probability that the conjunction is false, there are guaranteed to be some worlds at which it is false. Therefore, our historian will not be justified in believing the conjunction. However, they will also not be justified in believing its negation. The negation has a high probability. So, it only needs to be true throughout a very narrow sphere of worlds in order for the historian to be justified in believing it. However, we already know that it is false at all the most normal worlds consistent with our historian's evidence. So, despite its high probability, our historian is not justified in believing the negation of this conjunction.

Thus, we are able to secure the result that a high probability for p does not guarantee that p is justified. We are able to establish that a low probability for p does establish that p is not justified. And we are able to rule out justified inconsistent belief sets. Call the resulting theory the 'normic risk minimization' conception of justification. The normic risk minimization approach is both natural, and gives us an attractive picture of the relationship between justified belief and justified credence. However, it still bears clear similarities to the normic support approach. And the normic support approach faces other problems. In particular, it seems to allow for justified Moore paradoxical beliefs (Littlejohn and Dutant (2020)). In the final section I will explain how the normic risk minimization approach avoids this problem.

6 Justification, Knowledge, and Moore Sentences

Littlejohn and Dutant (2020) argue that the normic support approach allows for Moorean absurdities such as 'p and I don't know that p' to be justified. They present two arguments for this claim. We'll consider each in turn.

In what follows, read '$E \blacksquare\!\!\rightarrow p$' as '$E$ support's p', and Kp as 'S knows that p'. The standard normic support approach and the normic risk minimization approach both validate *reflexivity*:

$(E\&p) \blacksquare\!\!\rightarrow p$'

Littlejohn and Dutant observe that *Reflexivity* entails the following:

$(E\&p\&\neg Kp)\blacksquare\!\!\rightarrow p\&\neg Kp$

That is, if S's evidence contains E, p, and $\neg Kp$, then their total evidence supports the proposition $\ulcorner p\&\neg Kp\urcorner$. And, as Littlejohn & Dutant point out, there is nothing in the standard normic support account to prevent one's evidence from containing both p and $\neg Kp$. After all, p and $\neg Kp$ are consistent. So, why could they not both be true at the most normal worlds consistent with one's evidence? If this is right, then the normic support approach tells us that a subject's evidence could justify them in believing the Moorean absurdity 'p and I don't know that p'. Not only is this absurd, but it also undermines one of the primary motivations for denying that we are justified in believing lottery propositions: if we were justified in believing lottery propositions then we would be justified in believing 'my ticket will lose and I don't know that my ticket will lose'.

Littlejohn and Dutant note that we can avoid the possibility of one's evidence containing both p and $\neg Kp$ by holding that only knowledge can be evidence ($E = K$). I am not unsympathetic to $E = K$. If the only way to avoid Moorean absurdities was to endorse $E = K$ this is a price I'd be willing to pay. However, it is not clear that $E = K$ is required to rule out the possibility of justified Moore paradoxical beliefs on the normic risk minimization approach. It follows from three plausible assumptions that p and $\neg Kp$ cannot both be part of one's evidence.

The first assumption is that if E is S's evidence at w_1, and w_2 is a world that is both compatible with E and more normal than w_1, then E will be S's evidence set at w_2 as well. The basic motivation for this principle is the internist thought that two internal duplicates will always possess the same evidence. This is, of course, controversial. If $E = K$ then, even if w_1 and w_2 are subjectively indistinguishable, one could easily know more at w_2 than at w_1. After all, it is plausible that the more normal the world, the more of one's beliefs will constitute knowledge. However, if $E = K$ then we already have an explanation for the fact that p and $\neg Kp$ cannot both be part of one's evidence. So this is not something we should be worried about.

The second assumption is that one cannot know that p if there is a relevant risk of p being false. That is, if p is false at a sufficiently nearby world at which one believes that p in the same way, then one's belief does not constitute knowledge. I take this assumption to capture the basic motivation behind the standard safety condition on knowledge. However, unlike the standard safety condition it allows that one cannot know that p when p has a low probability conditional on one's evidence. I believe this is a welcome result. But, it won't be entirely uncontroversial as it does lead to failures of multi-premise closure.

The final assumption is that the relation of comparative normalcy of worlds is transitive. So, if w_1 is more normal than w_2, and w_2 is more normal than w_3, then w_1 is more normal than w_3. That the relation of comparative normalcy is transitive is often assumed (see for example Stalnaker (2005, 2015), Smith (2007, 2016), Greco (2014), and Goodman & Salow (2018)). However, it is not uncontroversial. Carter (2019) argues that what counts as normal is contingent, and this means that, just because something is normal at the actual world, doesn't mean that it is normal relative to some other world that itself counts as normal relative to the actual world.

This is not the place for a full discussion of whether normality iterates. I will simply observe that, following Smith (2007, 2010, 2016, and Forthcoming a), we are employing a very specific notion of normality tied closely to explanation. In particular, Smith has us thinking about normality in terms of idealization: worlds are ranked in terms of how idealized they are. It doesn't seem to be a contingent matter which worlds are more idealized than others (relative to the ceteris paribus laws governing our world at least). So, even if Carter's concerns do apply to the notion of normality employed by typical speakers, it is not clear that they will apply to the notion of normality at play here.

If these assumptions are correct then p and $\neg Kp$ cannot both be part of one's evidence. In order for both p and $\neg Kp$ to be part of one's evidence it would have to be the case that both p and $\neg Kp$ are true at all of the most normal worlds consistent with one's evidence. And this entails that Kp must be false at all of these worlds. There are only a few ways in which one can fail to know that p: one's belief could be false, unjustified, or at risk of being false given how it was formed. Since p is part of one's evidence, p cannot be false at any of the most normal worlds consistent with one's evidence. Moreover, by *Reflexivity*, one's belief that p is justified. So, there cannot be any risk of p being false at any of the most normal worlds consistent with one's evidence. So, the only way for $\neg Kp$ to be true at all the most normal worlds consistent with one's evidence would be for one's belief that p to be unjustified at all the most normal worlds consistent with one's evidence. Call the set of the most normal worlds consistent with one's evidence N_{E}. If E is your evidence, then E will also be your evidence at every world in N_{E}. Moreover, by transitivity of the 'more normal than' relation, the most normal E worlds relative to actuality will be identical to the most normal E worlds relative to every world in N_{E}. So, the most normal E worlds relative to any world in N_{E} will just be the worlds in N_{E}. But there is no risk of p being false at any of the worlds in N_{E}. So, it is true at every world in N_{E} that there is no risk of p being false at the most normal E worlds. Therefore, one's belief that p is also justified at every world in N_{E}. So, if one is justified in believing p, then one is in a position to know that p at all the most normal worlds consistent with one's evidence. As a result, p and $\neg Kp$ cannot both be part of one's evidence.

This blocks off the first route to Moorean absurdities. However, Littlejohn and Dutant identify a second route. They start by asking whether evidence for p is always evidence that one knows that p:

$$E\blacksquare \rightarrow p$$

$$\therefore\ E\blacksquare \rightarrow Kp$$

If evidence for E is always evidence for Kp then E will never be able to support $\ulcorner p \& \neg Kp \urcorner$. However, as Dutant and Littlejohn point out, this assumption is too strong. It plausibly fails in standard margin for error cases. Moreover, it can be seen to fail on the normic risk minimization conception of justification.[13]

However, if E can support p without supporting Kp this would seem to open the door to Moorean absurdities. Littlejohn and Dutant argue as follows. They start with *Rational Monotonicity*, which is validated by standard normic support theories of justification:

$$E\blacksquare\rightarrow p$$

$$\neg(E\blacksquare\rightarrow\neg q)$$

$$\therefore\ (E\&q)\blacksquare\rightarrow p$$

Since it is possible for E to support p without supporting Kp it will be possible to get cases in which $E\blacksquare\rightarrow p$ and $\neg(E\blacksquare\rightarrow\neg\neg Kp)$. Together with *rational monotonicity* this entails the following: $(E\&\neg Kp)\blacksquare\rightarrow p$. This is already an uncomfortable result. However, together with *reflexivity* we get $(E\&\neg Kp)\blacksquare\rightarrow\neg Kp$, and then by *agglomoration* we can derive the Moorean absurdity $(E\&\neg Kp)\blacksquare\rightarrow p\&\neg Kp$.

This argument breaks down in two places on the normic risk minimization conception of justification. Firstly, and unsurprisingly, the final step fails since the normic risk minimization approach invalidates agglomoration. However, it would still be an uncomfortable result if it was possible to derive $(E\&\neg Kp)\blacksquare\rightarrow p$. Thankfully this cannot be derived as *Rational Monotonicity* fails on the normic risk minimization approach. *Rational monotonicity* tells us that if our evidence supports p, and it does not support $\neg q$, then if we add q to our body of evidence, our new body of evidence will still support p. Now, suppose that our body of evidence supports p: p has a reasonably high probability conditional on our evidence, and it is true throughout a large enough sphere of normal worlds. Furthermore, suppose that our evidence supports neither q nor $\neg q$. More precisely, assume that $\neg q$ is true at all of the most normal worlds consistent with our evidence, but that, relative to some of these worlds there is a risk of q being true. So, there is some (perhaps fairly distant) world in the relevant

[13]In order for E to support p there must be no risk of p being false at any of the most normal worlds consistent with E. In order for E to support Kp there would have to be no risk of Kp being false at any of the most normal worlds consistent with E. We have seen that if E supports p then Kp will be true at all of the most normal worlds consistent with E. But that is consistent with there being a risk of Kp being false at some of these worlds: Suppose that w_1 is among the most normal worlds consistent with E. Suppose E supports p, so both p and Kp are true at w_1. Moreover, since there is no risk of $\neg p$ at w_1, p will be true throughout a sphere of worlds less normal than w_1. However, Kp needn't be true throughout this sphere of worlds. For Kp to be true throughout this sphere of worlds there would have to be no risk of p being false at any of the worlds in this sphere. Suppose w_2 is the least normal world that is sufficiently nearby to w_1 to affect whether there is a risk of $\neg p$ at w_1. And suppose that w_3 is the least normal world that is sufficiently nearby to w_2 to affect whether there is a relevant risk of $\neg p$ at w_2. Since w_3 will not be within the sphere of worlds relevant to determining whether there is a relevant risk of $\neg p$ at w_1, it is compatible with E supporting p, and p being known at w_1 that p is false at w_3. And if this is the case, there will be a risk of p being false at w_2, meaning there is a risk of p not being known at w_1. So, whilst E will support p, it will not support Kp.

sphere at which q is true. Furthermore, suppose that $P(p|q)$ is extremely low. If this is the case then, in order for $\ulcorner E \,\&\, q \urcorner$ to support p, p would have to be true throughout an extremely wide sphere of worlds. It could easily fail to be true throughout such a sphere despite being true throughout the (potentially far narrower) sphere of worlds necessary for there to be no risk of $\neg p$ at the most normal worlds consistent with E. In cases like this, we will get failures of *rational monotonicity*. So, the derivation of Moorean absurdities is blocked.

7 Conclusion

To conclude, the normic risk minimization approach to justification is natural, and gives us an appealing account of the relationship between justified belief and justified credence: if one is justified in having a high credence in p then this means that one's evidence will only have to support p throughout a fairly narrow sphere of normal worlds in order to be justified. So, typically, but not always, a justified high credence in p will bring with it a justified belief in p. Moreover, when the probability of p is low there is guaranteed to be a risk of p. So, a justified low credence in p will always render belief in p irrational. This approach is able to give satisfactory solutions to both the lottery and preface paradoxes. It has affinities to knowledge first approaches to justification since, on certain reasonable assumptions, it entails that one is justified in believing that p only if one would be in a position to know that p in normal circumstances. However, it does not analyze justification in terms of knowledge. And, unlike the normic support approach upon which it is based, it does not allow the generation of Moorean absurdities.

References

[1] Adler, J. 2002. *Belief's own ethics.* Cambridge: MIT University Press.

[2] Backes, M. 2019a. Normalcy, Justification, and the Easy Defeat Problem. *Philosophical Studies* 176 (11). 2877-2895.

[3] Backes, M. 2019b. A Bitter Pill for Closure. *Synthese* 196 (9). 3773-3787.

[4] Bear, A. & Knobe, J. 2017. Normality: Part descriptive, part prescriptive. *Cognition* 167. 25-37.

[5] Blome-Tillman. 2015. Sensitivity, Causality, and Statistical Evidence in Courts of Law. *Thought: A Journal of Philosophy* 4. 102-112.

[6] Carter, S. 2019. Higher Order Ignorance Inside the Margins. *Philosophical Studies* 176. 1789-1806.

[7] Ebert, P, Smith, M. & Durbach, I. 2020. Varieties of Risk. *Philosophy and Phenomenological Research* 101 (2). 432-455.

[8] Evnine, S. J. 1999. Believing conjunctions. *Synthese* 118. 201–227.

[9] Foley, R. 1992. The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly* 29. 111–121.

[10] Goodman, J. & Salow, B. 2018. Taking a Chance on KK. *Philosophical Studies* 175. 183-196.

[11] Greco, D. 2014. Could KK be OK? *The Journal of Philosophy* 111. 169-197.

[12] Ichikawa, J. 2014. Justification is Potential Knowledge. *Canadian Journal of Philosophy*, 44. 184-206.

[13] Leitgeb, H. 2017. *The stability of belief.* Oxford: Oxford University Press.

[14] Littlejohn, C. 2012. Lotteries, probabilities, and permissions. *Logos and Episteme*, 3, 509–514.

[15] Littlejohn, C., & Dutant, J. 2020. Justification, Knowledge, and Normality. *Philosophical Studies.* 177 (1593-1609).

[16] Loets, A. Forthcoming. Choice Points for a Theory of Normality. *Mind*.

[17] Makinson, D. C. 1965. The paradox of the Preface. *Analysis* 25. 205-207.

[18] Nelkin, D. 2000. The lottery paradox, knowledge and rationality. *Philosophical Review* 109. 373– 409.

[19] Peet, A. A Risk Analysis.

[20] Praolini, F. 2019. No Justificatory Closure without Truth. *Australasian Journal of Philosophy* 97 (4). 715-726.

[21] Pritchard, D. 2015. Risk. *Metaphilosophy* 46 (3). 436-461.

[22] Rosenkranz, S. 2018. The Structure of Justification. *Mind* 127. 309-338.

[23] Ryan, S. 1991. The preface paradox. *Philosophical Studies* 64. 293–307.

[24] Smith, M. 2007. Ceteris Paribus Conditionals and Comparative Normalcy. *Journal of Philosophical Logic* 36. 97-121.

[25] Smith, M. 2010. What Else Could Justification Be? *Noûs* 44 (1). 10-31.

[26] Smith, M. 2016. *Between Probability and Certainty* Oxford: Oxford University Press.

[27] Smith, M. Forthcoming a. The Hardest Paradox for Closure. *Erkenntnis*

[28] Smith, M. Forthcoming b. Risky Belief. *Philosophy and Phenomenological Research.*

[29] Stalnaker, R. 2005. On Logics of Knowledge and Belief. *Philosophical Studies* 128. 169-199.

[30] Stalnaker, R. 2015. Luminosity and the KK Thesis. In Goldberg, S. (ed) *Externalism, Self Knowledge, and Skepticism: New Essays.* Cambridge: Cambridge University Press.

[31] Sutton, J. 2007. *Without Justification* Oxford: Oxford University Press.

[32] Williamson, T. 2009. Probability and Danger. *The Amherst Lecture in Philosophy.* 4. 1-35.

[33] Williamson, T. 2011. Improbable Knowing. In Dougherty, T. (ed) *Evidentialism and its Discontents* Oxford University Press.

[34] Williamson, T. 2014. Very Improbable Knowing. *Erkenntnis* 79 (5). 971-999.