

The Hardest Paradox for Closure

Martin Smith

According to the principle of *Conjunction Closure*, if one has justification for believing each of a set of propositions, one has justification for believing their conjunction. The lottery and preface paradoxes can both be seen as posing challenges for Closure, but leave open familiar strategies for preserving the principle. While this is all relatively well-trodden ground, a new Closure-challenging paradox has recently emerged, in two somewhat different forms, due to Marvin Backes (2019a) and Francesco Praolini (2019). This paradox synthesises elements of the lottery and the preface and is designed to close off the familiar Closure-preserving strategies. By appealing to a normic theory of justification, I will defend Closure in the face of this new paradox. Along the way I will draw more general conclusions about justification, normalcy and defeat, which bear upon what Backes (2019b) has dubbed the ‘easy defeat’ problem for the normic theory.

I BACKGROUND: THE LOTTERY AND PREFACE PARADOXES

Consider the following principle:

If one has justification for believing each of P_1, P_2, \dots, P_n , then one has justification for believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$.

According to this principle, the set of propositions that one has justification for believing is closed under the operation of taking conjunctions – we might call it *Conjunction Closure* or, simply, *Closure*. In one way, the principle seems difficult to deny. Suppose I endeavour to believe all and only those propositions for which I have justification. If Closure fails, there will be possible situations in which I ought to believe each of a series of propositions while refraining from believing their conjunction. And yet, it’s unclear how I would actually *go about* following such a recommendation, even if I accepted it. It’s unclear that there is any psychological difference between believing P_1 , believing P_2 , ... believing P_n and believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$ (see Evidentiary, 1999, section 7, Douven, 2002, pp395). If a person has expressed belief in P_1 , and in P_2 , ... and in P_n we would, without a second thought, describe the person as believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$.

Suppose I’m filling in a large truth table. If I put a ‘T’ in the P_1 column and a ‘T’ in the P_2 column ... and a ‘T’ in the P_n column, you would describe me as believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$. Even if I hesitated when it came to the final $P_1 \wedge P_2 \wedge \dots \wedge P_n$ column, this wouldn’t obviously make any difference – it would merely look as though I don’t know how to complete a truth table for ‘ \wedge ’. One might claim that it is, at least, possible to *assert* each of a series of propositions without asserting their conjunction. But even this is unclear. If I assert, in sequence, P_1, P_2, \dots, P_n it would be natural for you to describe me as having asserted $P_1 \wedge P_2 \wedge \dots \wedge P_n$. One could insist that, in order to count as asserting the latter, I must literally voice the ‘and’s – but our ordinary practice of reporting assertions doesn’t appear to support any such requirement.

These kinds of points have I think been somewhat neglected in discussions of Closure¹ – but I won't attempt to expand upon them here. What discussions of Closure *have* tended to focus on are certain problems that the principle appears to generate. Suppose I hold a single ticket – ticket #72 – in a fair 100-ticket lottery with a single winner. Suppose I know that the lottery has been drawn and I'm yet to hear the result, but I already believe that ticket #72 has lost on the grounds that there are 99 losing tickets and only one winner. It's plausible that I have justification for believing this – given my evidence, it is highly likely to be true. But if I have justification for believing (P_{72}) that ticket #72 has lost, I must also have justification for believing (P_1) that ticket #1 has lost, and for believing (P_2) that ticket #2 has lost, and so on up to P_{100} . After all, my evidence for each of these propositions is exactly the same. By stipulation, I know that one of the tickets has won, in which case I also have justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$. If Closure holds, I must have justification for believing an outright contradiction – $P_1 \wedge P_2 \wedge \dots \wedge P_{100} \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100})$.

This is a version of the lottery paradox (Kyburg, 1961). The only way to preserve Closure, in the face of this paradox, is to insist that I lack justification for believing some of $P_1, P_2 \dots P_{100}$. Since I have the same evidence for each of these propositions, given a suitable assumption to the effect that justification is a function of evidence, it will follow that I lack justification for believing each of $P_1, P_2 \dots P_{100}$ ². Perhaps the fact that there is one winning ticket and 99 losers gives me justification for believing, of any given ticket, that it is *very likely* to have lost, but does not give me justification for believing that it *has* lost.

Suppose I have written a book in which I make 100 logically independent factual claims $P_1, P_2 \dots P_{100}$. Suppose each claim has been meticulously checked and it is stipulated that I've secured justification for believing each one. I'm well aware however that I'm fallible and that comparably ambitious books have always turned out to contain falsehoods in the past. As a result, I feel compelled to write in the preface that there are bound to be some falsehoods in the book. Plausibly I have justification for believing this – given my evidence, it's highly likely to be true. But if I have justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ and Closure holds, I must have justification for believing an outright contradiction – $P_1 \wedge P_2 \wedge \dots \wedge P_{100} \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100})$.

This is a version of the preface paradox (Makinson, 1965). The only way to preserve Closure, in the face of this paradox, is to insist that I lack justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee$

¹ An exception is Evident (1999). One reason, perhaps, for this relative neglect is that the transition from the n propositions $P_1, P_2, \dots P_n$ to the single proposition $P_1 \wedge P_2 \wedge \dots \wedge P_n$ would seem to involve an *inference* – an instance of Conjunction Introduction – and surely it should be possible, in principle, to hold back from drawing an inference. Such a thought is too quick, however. As some theorists have observed, the term 'inference' tends to be used in two rather different ways in philosophy (Burgess, 1981, pp103, Harman, 1986, pp3-4): Sometimes it is used to mean a kind of psychological process whereby one acquires new beliefs, and sometimes it is used to mean an ordered pair consisting of a sequence of premises and a conclusion (or perhaps multiple conclusions). But there is no guarantee whatsoever that an inference in the latter 'logical' sense will always correspond to an inference in the former 'psychological' sense.

² Some have proposed a view on which one can justifiably believe, of *some* lottery tickets, that they have lost, but one cannot justifiably believe this about *too many* of the tickets simultaneously (Harman, 1986, pp70-71, Kroedel, 2012, see also Douven, 2008). This allows us to retain a certain closure principle on which, if one has justification for simultaneously believing P_1 and believing $P_2 \dots$ and believing P_n , then one has justification for believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$. The view has significant costs however – whether one is justified in believing, of a given lottery ticket, that it has lost will no longer be purely a function of the evidence on which the belief is based, but will depend upon how many similar beliefs one has already formed. I won't discuss this view here. One further criticism of this proposed solution to the lottery paradox is that it may not extend to the preface paradox (see Eder, 2015).

$\sim P_{100}$. Perhaps the fact that I'm fallible and that carefully checked claims occasionally turn out to be false gives me justification for believing that it is *very likely* that the book contains falsehoods, but does not give me justification for believing that the book *does* contain falsehoods. Even if we pursue this route, however, the preface paradox has a sting in the tail; if Closure is correct then not only will I *lack* justification for believing that $P_1 \wedge P_2 \wedge \dots \wedge P_{100}$ is false, I will *have* justification for believing that $P_1 \wedge P_2 \wedge \dots \wedge P_{100}$ is *true* – justification for believing that the book is completely falsehood-free. This is a serious cost, as many have rightly emphasised (Christensen, 2004, chap. 3). Whether accepting this is more costly than abandoning Closure is I think less clear – but a topic for another occasion.

These two familiar paradoxes are, in a way, opposite sides of the same coin. In each case we have a logically inconsistent set of 101 propositions: $\{P_1, P_2, \dots, P_{100}, (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100})\}$. In the lottery paradox, it is stipulated that I have justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ (justification for believing that some ticket has won) and it is assumed that I must have justification for believing each of P_1, P_2, \dots, P_{100} (justification for believing, of each ticket, that it has lost). In the preface paradox, it is stipulated that I have justification for believing each of P_1, P_2, \dots, P_{100} (justification for believing each claim in the book) and it is assumed that I must have justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ (justification for believing that some of the claims in the book are false). These stipulations should be acceptable to all non-sceptics, but the assumptions provide an avenue of response for the defenders of Closure.

Put differently: In the lottery paradox, $\sim(P_1 \wedge P_2 \wedge \dots \wedge P_{100})$ is meant to be certain, but the evidence that I have for believing each of P_1, P_2, \dots, P_{100} – the evidence that there is one winning ticket and 99 losers – leaves each of these propositions with a distinctive kind of uncertainty. By denying that belief can be justified on the basis of such 'purely statistical' evidence one can avoid the paradox and preserve Closure (see, for instance, Ryan, 1996, Nelkin, 2000, Smith, 2010, 2016, section 3.1, Smithies, 2012). In the preface paradox, it is effectively left open what kind of evidence I have for believing each of P_1, P_2, \dots, P_{100} , allowing us to fill in the details as we see fit, but the evidence I have for believing $(\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100})$ is fixed – it consists in the fact that I am fallible and that comparably ambitious books have always turned out to contain falsehoods. By denying that belief can be justified on the basis of this 'pessimistic inductive' evidence, one can avoid the paradox and preserve Closure (see, for instance, Ryan, 1991, Kaplan, 2013, Kim, 2015, Smith, 2016, section 4.1).

In each paradox, a substantial assumption about epistemic justification is needed – it isn't possible to disprove Closure using stipulations alone. Or is it? In recent work, a new Closure-challenging paradox has been set out, in somewhat different forms, by Marvin Backes (2019a) and Francesco Praolini (2019). This paradox, which combines elements of the lottery and the preface, relies on no obvious assumptions about epistemic justification and appears to resist the familiar Closure-preserving strategies.

II THE HYBRID PARADOX

Following Praolini, suppose again that I have secured justification for believing each of 100 logically independent factual claims $P_1, P_2 \dots P_{100}$ which I compile in a book. Suppose I then send my book manuscript to the Perfectly Omniscient Press for consideration for publication, and am informed by their perfectly omniscient and truthful referee that there is one false claim in the book, providing me with justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$. If Closure holds, I must have justification for

believing an outright contradiction – $P_1 \wedge P_2 \wedge \dots \wedge P_{100} \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100})$. As in the preface paradox, it is simply stipulated that I have justification for believing each of $P_1, P_2 \dots P_{100}$ – my evidence in favour of these propositions is left unspecified and can be filled in as we wish. As in the lottery paradox, it is simply stipulated that I have justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ – this proposition is meant to be certain. As a result, this ‘hybrid’ paradox does not appear to rely on any substantial assumptions³.

In fact, $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ doesn’t quite exhaust the content of the referee’s report, as Praolini describes it. In Praolini’s example, the referee reveals not just that there is *at least* one falsehood in the book, but that there is *exactly* one falsehood in the book – which entails, but is not entailed by, $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$. The extra content of the report could be captured by making this an *exclusive* rather than an inclusive disjunction⁴. The extra content, however, is inessential to the paradox, and makes no difference for what follows – with one possible exception which I will note.

As Praolini points out, when I receive the report from the perfectly omniscient referee, one might think that this serves to *defeat* my justification for believing the claims in the book. In this case, when I acquire justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$, I would lose justification for believing $P_1, P_2 \dots P_{100}$ and there would be no single time at which I have justification for believing all 101 propositions. While this is one possible way of blocking the paradox, according to Praolini it is implausible to think that the referee’s report would defeat my justification for everything that I’ve written in the book. Rather than imagining a book with a mere 100 claims, Praolini considers a book which contains ‘all and only logically independent propositions that you are justified to believe’ (Praolini, 2019, pp720). For any set of propositions, there may be many subsets with mutually logically independent members, and none of these need be maximal – but suppose we imagine a book detailing some suitably large set of logically independent propositions that I believe. Now the stakes are even higher. Defeating all of the claims in the book would be tantamount to defeating my justified beliefs *en masse* – and surely the referee’s report would not have this effect.

Backes (2019a) describes a case that is structurally similar: Suppose I am slipped a pill that ensures that some small proportion of the justified beliefs that I form, during a certain period, will be false. When the period has elapsed, I learn of the pill and its effects, putting me in much the same situation as in Praolini’s example. Backes is also aware that Closure could be saved if the information

³ There are some anticipations of this paradox in earlier literature. For instance, the paradox bears a relation to Ryan’s ‘third version’ of the preface paradox (Ryan, 1991, pp304) in which one receives a defeasible report to the effect that one’s book contains a falsehood. She concludes that, if one has genuinely secured justification for believing each claim in the book, then one would not be justified in believing the report and should simply dismiss it. This closure-preserving strategy is not available in Praolini’s case. In the ‘homogeneous preface paradox’ described by Easwaran and Fitelson (2015, section 3) one writes a well-researched ambitious factual book arguing precisely that all well-researched ambitious factual books contain at least one error. In this case, Easwaran and Fitelson suggest that one has something much stronger than the usual ‘pessimistic inductive’ evidence to the effect that one’s book contains an error – which, if correct, would give the paradox something like the same character as Praolini’s. In fact, I think it’s far from clear that this suggestion is correct – something I will return to in n9.

⁴ If we simply replace the ‘ \vee ’s with binary ‘exclusive-or’s then this won’t give us the desired content (no matter where we add brackets). We could instead make use of an n-place exclusive-or function which returns the value true iff exactly one of its arguments is true. Thanks to Jim Pryor here.

about the pill is taken to defeat all of the justified beliefs that I formed during this period – but he too regards it as implausible that this information could have such a devastating defeating effect⁵.

For Backes, the lesson of this paradox is that Closure fails. For Praolini, the lesson of the paradox is that either Closure fails or justification is *factive* – one can only have justification for believing true propositions (see for instance Sutton, 2007, Littlejohn, 2012, Williamson, forthcoming). As Praolini points out, if justification were factive, then the very set-up of the paradox would be impossible. If I had justification for believing each of $P_1, P_2 \dots P_{100}$, and justification were factive, then the omniscient, truthful referee couldn't inform me of $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ – for this would be false.

In the next section, I will set out a rather different response to the paradox – one that revives the idea that the referee's report may function as a kind of *defeater*. The first thing to observe is that, in order to block the paradox, it is not necessary that the report defeat *all* of the claims in the book – it is enough that it defeat *some*. Since the report is so general, and fails to single out any individual claims, the idea that it would serve to defeat some claims and not others may look like a nonstarter. I will argue that this idea is more promising than it first appears.

III THE PRINCIPLE OF DIFFERENTIAL DEFEAT

Suppose I've been invited to a drinks reception, and I know that a very eminent, world-leading primatologist will be in attendance. In preparation for the event, I arm myself with three 'primate facts' that I can causally drop into the conversation in case the primatologist and I are introduced. First, I read in the current edition of *Encyclopedia Britannica* that bonobos are capable of passing the mirror self-recognition test (P_1). Second, I read in a newspaper article that Madagascar was once home to lemurs that were larger than humans (P_2). Finally, a few days before the reception, I hear in conversation that the barbary macaque is the only species of old world monkey that lacks a tail (P_3). I come to believe each of P_1, P_2 and P_3 and, plausibly, I am justified in doing so. At the reception I am introduced to the primatologist and, over-eager to impress, I blurt out all three 'facts' in quick succession. The primatologist furrows her brow and says 'I'm afraid something that you just said there is wrong'. Before she can elaborate further, however, she is quickly whisked away to meet another guest.

How, at this point, should I revise my beliefs? If I accept what the primatologist says, it's clear that I can't just continue to believe each of P_1, P_2 and P_3 – that would be irrational. One thing I could do is to give up all of these beliefs, and suspend judgment on P_1, P_2 and P_3 . This may be a permissible response to my new evidence – but, in a way, it seems like an overreaction. After all, my evidence for these three propositions is not equal – it's natural to think that my justification for believing P_1 is *stronger* than my justification for believing P_2 which, in turn, is *stronger* than my

⁵ Dutant and Littlejohn (2020, section 2) discuss two further examples which have the same structure as those described by Praolini and Backes. First, imagine a person undergoing an eye test who is asked to identify various letters and numbers on a series of slides. While she is able to answer all of the questions easily and forms a series of justified perceptual beliefs over the course of the test, she is subsequently told by the optometrist that she made one error. Second, imagine a judge who has convicted many defendants over her career. Suppose that, in each such case, there was strong incriminating evidence and she justifiably believed the defendant to be guilty. Suppose she is then reliably informed that one of the people she has convicted is innocent. Dutant and Littlejohn also consider, and dismiss, the possibility that the new evidence from the optometrist or the informant could defeat all of the beliefs that fall within its scope – all of the eye-test beliefs or guilt beliefs. Though I will focus on Praolini's example in the main text, I will have a bit more to say about these examples in n7 – and I hope to discuss them at length elsewhere.

justification for believing P_3 . While the current edition of Encyclopedia Britannica is a highly reliable source for information about primates, a newspaper is somewhat less reliable and a snippet of conversation is less reliable still. Another permissible response to the new evidence, I suggest, is to retain my beliefs in P_1 and P_2 and to give up my belief in P_3 . If this suggestion is right, then it is only my justification for P_3 that is defeated by the primatologist's remark – my justification for P_1 and P_2 survives.

Turning back to Praolini's example, if we imagine a book containing most of the propositions that I justifiably believe, this will include claims such as 'Two and two is four' and 'I am not a turnip' through to claims such as 'Edinburgh Waverly has the highest annual footfall of any train station in Scotland' and 'Oswald acted alone'. When I discover that the book contains a falsehood, it's highly implausible that all of these claims suddenly become *equally doubtful*. Rather, the claims that are made most doubtful by this discovery are the ones that *were the most doubtful to begin with*.

Suppose one has justification for believing each proposition in a set $\{P_1, P_2, \dots, P_n\}$. According to what I will call the *Principle of Uniform Defeat*, if one learns $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ then this will defeat one's justification for every proposition in $\{P_1, P_2, \dots, P_n\}$. The Principle of Uniform Defeat does have an initial appeal and may be a consequence of principles that some philosophers have explicitly endorsed – such as Ryan's 'avoid falsity principle' (Ryan, 1996)⁶. I propose instead a *Principle of Differential Defeat*: If one learns $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ then this will defeat one's justification for all and only those propositions in $\{P_1, P_2, \dots, P_n\}$ that were the least justified, prior to the discovery.

Suppose we rank the claims in the book according to how justified they are:

.
 .
 .
 P₇₆, P₁₂, P₈₂
 P₁₉, P₇₁, P₄, P₉₆
 P₅₀, P₃₁
 P₆₄, P₃₉, P₇, P₉₀

According to the Principle of Differential Defeat, when I learn $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ this serves to defeat my justification for believing all and only those claims in the bottom tier. In this case, when I receive the referee's report, I lose justification for believing P_{64} , P_{39} , P_7 and P_{90} and the paradox is blocked.

The Principle of Differential Defeat offers a way of resolving the hybrid paradox without abandoning Closure or embracing the claim that justification is factive. Praolini is aware of this potential solution to the paradox and offers four replies (see Praolini, 2019, pp722-723), which I will consider in turn. First, Praolini suggests that, far from defeating some of the claims in the book, the referee's report should actually *increase* my justification for each of the claims. Recall that, in Praolini's example, the referee informs me that there is *exactly* one falsehood in the book and no more – and, according to Praolini, this should come as *good news*. To motivate this, he reiterates

⁶ According to the Avoid Falsity Principle, for any competing set of statements L , if one has good reason to believe every member of L , and one also has good reason to believe that at least one member of L is false or is justified in suspending judgment about whether at least one member of L is false, then one is not epistemically justified in believing any of the members of L (Ryan, 1996, pp130). Given Ryan's definition of what it is for a set of statements L to be 'competing', this appears to be automatically satisfied in case one has good reason for believing that at least one member of L is false.

some reasoning that is familiar from the preface paradox; given that I'm fallible, and there are so many claims in the book, we should expect it to contain *numerous* falsehoods. As a result, to discover that there is only *one* false claim in the book is to discover that things have turned out much better than expected, and the discovery should boost my justification for each individual claim. As discussed above, this reasoning does have a certain appeal – but it is reasoning that a defender of Closure is already committed to rejecting. If Closure holds then, prior to receiving the referee's report, I have justification for believing that there are no false claims in the book. As a result, there is at least one sense in which I should not 'expect' the book to contain numerous falsehoods, and in which the report does not represent better news than expected. Praolini's first reply should, then, leave a defender of Closure unmoved – it effectively takes it for granted that Closure fails.

Second, Praolini suggests that, since the referee report fails to specify any particular claims, it is counterintuitive that it would defeat some claims and not others. That is, it is counterintuitive that the referee report should defeat only P_{64} , P_{39} , P_7 and P_{90} – even if these do happen to be the least justified in the book. As noted above, there is something attractive about this suggestion, but it fails to stand up to scrutiny. Even if a defeater weighs equally against each of a set of justified beliefs, the beliefs themselves will typically vary with respect to how *vulnerable* they are to defeat – as the above examples illustrate. Broadly speaking, the stronger one's justification for a proposition, the more resistant it is to defeat. The reason that P_{64} , P_{39} , P_7 and P_{90} buckle under the strain of the report is not that it weighs extra heavily against them (it doesn't) – it is because these are the claims that are least able to bear the weight.

Third, according to Praolini, it is ad hoc to maintain that the referee's report only defeats the least justified claims in the book, when a conjunction of other claims in the book may have an even lower level of justification. That is, it would be ad hoc to maintain that the referee's report defeats the least justified claims P_{64} , P_{39} , P_7 , P_{90} if a conjunction of further claims – say, $P_{19} \wedge P_{71} \wedge P_4 \wedge P_{96}$ – was less justified still. I am inclined to think that Praolini's third reply, like the first, effectively begs the question against those who would defend Closure.

Consider the following principle which we might call *Comparative Closure*:

One's justification for believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is no weaker than one's justifications for believing each of P_1, P_2, \dots, P_n .

Comparative Closure is a stronger principle than Closure, but the most common motivations for accepting the latter would seem to carry over to the former. This would certainly seem to be so for the motivation that I sketched at the outset; if one automatically counts as believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$ whenever one believes P_1 , believes P_2 , ... and believes P_n then this gives us reason to deny, quite generally, that one's epistemic standing with respect to $P_1 \wedge P_2 \wedge \dots \wedge P_n$ could be worse than one's epistemic standing with respect to each of P_1, P_2, \dots, P_n . In any case, defenders of Closure are under strong pressure to accept Comparative Closure, which effectively rules out the kind of possibility that Praolini envisages.

To pursue this line of thought a little further, it is also very plausible that my justification for $P_1 \wedge P_2 \wedge \dots \wedge P_n$ cannot be any *stronger* than my justification for the least justified member of $\{P_1, P_2, \dots, P_n\}$. If we combine this with Comparative Closure, we can derive the following:

The degree of one's justification for believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is equal to the degree of one's justification for believing the least justified member $\{P_1, P_2, \dots, P_n\}$.

We might call this the *Minimum Conjunct Rule*. Using this rule, we could add conjunctions to the above ranking as follows:

$$\begin{array}{c} \cdot \\ \cdot \\ \cdot \\ P_{76}, P_{12}, P_{82} \dots P_{76} \wedge P_{12} \wedge P_{82} \\ P_{19}, P_{71}, P_4, P_{96} \dots P_{19} \wedge P_{71} \wedge P_4 \wedge P_{96} \\ P_{50}, P_{31} \dots P_{50} \wedge P_{31} \\ P_{64}, P_{39}, P_7, P_{90} \dots P_{64} \wedge P_{39} \wedge P_7 \wedge P_{90} \end{array}$$

I will return to this in the next section.

Finally, Praolini points out that, if the referee report defeats the least justified claims in the book, and the claims all happened to be *equally justified*, then they would all be defeated. This result is, indeed, unavoidable – if the propositions in a set are equally justified, then the Principle of Differential Defeat and the Principle of Uniform Defeat will make exactly the same predictions. Such a case would, however, be very different from what Praolini initially asks us to imagine. Suppose I'm looking at a row of cereal boxes on a supermarket shelf. Presumably I'm justified in believing, of each box, that it contains cereal. Suppose I then learn that one of the boxes is empty and has been placed on the shelf by mistake. This appears to be a case in which I learn that one amongst a set of equally justified propositions is false. After all, I have no more reason to think that any one box contains cereal than any other – I can't tell this just by looking. In this case, though, it is plausible that the new evidence *would* serve to defeat my justification for believing each proposition in the set. That is, it's plausible that I would no longer be justified in believing, of any one box, that it contains cereal – not without picking it up or looking inside. Even if we find the idea of en masse defeat implausible when it comes to Praolini's original example, we should not assume that this intuition will persist when the example has been adjusted in such a way as to ensure that all of the beliefs in question are equally justified⁷.

In any event, defenders of Closure may have a particular reason for tolerating en masse defeat in cases of this kind. If I have equal justification for each of a series of claims and I learn that exactly one of them is false, then the situation that I confront is very similar to that presented by the lottery paradox. Both situations involve a large set of propositions, each of which is very likely to be true but one of which is sure to be false. In both situations, the propositions are on a par, in that any one could be the false proposition just as easily as any other. Indeed, if we wrote each of the claims down on

⁷ Dutant and Littlejohn's eye test example may be another case in which one learns that one amongst a set of equally justified beliefs is false. The way that Dutant and Littlejohn describe it, the beliefs that the examinee forms about the letters and numbers on the eyecharts are equally sound from her own perspective, and she has no reason to regard any one as more doubtful than any other. If these beliefs are equally justified, the Principle of Differential Defeat predicts that they will all be defeated by the optometrist's revelation. (Unless the optometrist's words are themselves thought to be open to doubt, in which case the examinee faces a more complex epistemic situation akin to Ryan's third version of the preface paradox discussed in n3. I return to this kind of situation in n16). It is worth noting that Dutant and Littlejohn's judge example is quite different in this regard. One could not simply stipulate that the judge's guilt beliefs are all *equally justified* without considerably altering the case and making it far more contrived than what one would naturally imagine. In fact, even in the eye test example, the stipulation of equal justification may not be altogether natural – see <https://www.ismp.org/resources/misidentification-alphanumeric-symbols>. Thanks here to Philip Ebert.

slips of paper, then there would be one ‘winning ticket’, which featured the one false claim, and a multitude of ‘losing tickets’, which featured the true claims. To believe any particular claim would be tantamount to believing, of one particular ticket, that it’s a loser. For a defender of Closure, accustomed to denying that one can justifiably believe, of a single ticket, that it has lost a fair lottery, embracing en masse defeat in this situation is, I think, a small step. I will have a bit more to say about these sorts of cases in the final section⁸.

In this section I have provided a prima facie motivation for the Principle of Differential Defeat, and argued that this principle will allow us to preserve Closure in the face of the hybrid paradox. It remains to be shown that there is a viable theory of justification that will vindicate both Closure and the Principle of Differential Defeat. I turn to this next.

IV THE NORMIC THEORY

In both the lottery and the preface paradoxes, we are invited to infer that I have justification for believing a proposition from the premise that it is *highly likely*, given my evidence. In the lottery paradox the proposition in question is that ticket #72 has lost, while in the preface paradox the proposition in question is that the book contains falsehoods. The inference is a very tempting one. After all, most epistemologists are *fallibilists* who agree that one can have justification for believing a proposition even if one’s evidence doesn’t make it completely certain. But what else can we require, then, except that one’s evidence make the proposition *likely*? What else could the evidence *do*?

On reflection, I think that there *is* something else that the evidence might do. Sometimes our evidence in favour of a proposition P is such as to make the falsity of P *abnormal* in the sense of requiring special explanation. Suppose I wander into a room I’ve never been in before and notice that the wall before me appears to be red. Clearly this evidence makes it very likely that the wall before me really is red – but this is not its only effect. If the wall appears to me to be red, but it *isn’t* red, then there would have to be some explanation as to how this came to be – I’m undergoing a colour hallucination, the wall is illuminated by hidden red lights, I’ve suddenly been struck by colour blindness etc. Whatever the case, there has to be more to the story – it can’t ‘just so happen’ that the wall appears to me to be red, but isn’t red.

In contrast, the fact that there are 99 losing tickets and only one winner doesn’t generate the need for special explanation in the event that ticket #72 is the winner. If ticket #72 were to win, then I may be surprised and delighted (it is my ticket after all) – but I wouldn’t seek some special explanation as to how this could possibly have happened. Some ticket has to win the lottery and it might just as well be ticket #72 as any other. Although it would be very unlikely, there is a sense in which there would be nothing *abnormal* about this ticket being the winner (Vogel, 1999).

More controversially, the fact that I’m fallible and that comparably ambitious books have always contained falsehoods in the past does not generate the need for a special explanation in the

⁸ Another kind of situation which might be thought problematic for the Principle of Differential Defeat is one in which all of the propositions in a set are equally justified save for one, which has a higher degree of justification than the others. If I learn that some member of the set is false, the Principle of Differential Defeat predicts that this one proposition will survive, while the remainder are all defeated. It might seem odd, though, that I should retain justification for believing this proposition alone – particularly if my initial justification for believing it was only slightly higher than my justification for believing the others. This kind of example raises questions about the *granularity* of justification – questions about how finely we can differentiate different levels of justification. I will return to this in n13 and n17.

event that my book turns out to be falsehood-free. Once again, if this were the case then I may be surprised and delighted – but I wouldn't demand an explanation as to how this could possibly have happened. Recall that every claim in the book has been thoroughly researched and checked – and it could just turn out that my research has delivered the right result every time. Why shouldn't it? Although it would be very unlikely for every claim in the book to be true there is, once again, a sense in which there would be nothing *abnormal* about this turn of events.

Sometimes when we describe a situation as 'normal' or 'abnormal' we are simply making a claim about frequencies – a normal situation is one that frequently arises, while an abnormal situation is one that is infrequent. If this is our understanding of normalcy, then we *should* say that it would be abnormal for ticket #72 to win or for my book to be error free. This is not the only way that we use these terms however. If the lights in my house suddenly start to flicker, or my car fails to start when I turn my key in the ignition and I remark 'that's not normal', I'm not just pointing out that this is something rare or infrequent – part of what I'm saying is precisely that there needs to be some special explanation for what is occurring.

Say that evidence E normically supports a proposition P just in case, given E, the situation in which P is false would be abnormal in the sense of requiring special explanation (Smith, 2010, 2016, 2018). The evidence that the wall appears to be red normically supports the proposition that the wall is red. The evidence that there are 99 losing tickets and one winning ticket does not normically support the proposition that my ticket has lost. The fact that I'm fallible and that comparable books have always turned out to contain falsehoods in the past does not normically support the proposition that my book contains a falsehood.

According to the normic theory of justification, one has justification for believing a proposition P just in case one's evidence normically supports P. In the lottery paradox, the normic theory predicts that I lack justification for believing, of any ticket, that it has lost the lottery – I lack justification for believing any of P_1, P_2, \dots, P_{100} . In the preface paradox, the normic theory predicts that I lack justification for believing that the book contains a falsehood – I lack justification for believing $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ ⁹. Not only does the normic theory offer a way of preserving Closure in the face of the lottery and preface paradoxes – it would appear to deliver a general validation of the principle: Suppose one has justification for believing P_1 , justification for believing P_2, \dots , justification for believing P_n . According to the normic theory, given one's evidence E, there would have to be a special explanation if P_1 were false and there would have to be a special explanation if P_2 were false ... and there would have to be a special explanation if P_n were false. What about $P_1 \wedge P_2 \wedge \dots \wedge P_n$? If $P_1 \wedge P_2 \wedge \dots \wedge P_n$ were false, then at least *one* of P_1, P_2, \dots, P_n would have to be false. Therefore, given one's evidence E, there would have to be a special explanation if $P_1 \wedge P_2 \wedge \dots \wedge P_n$ were false and, according to the normic theory, one has justification for believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$ (Smith, 2018, pp3870). The claim that the normic theory validates Closure will also permit of a more rigorous proof – given a certain formal development of the notion of normic support – which I will outline in the next section.

⁹ In the 'homogeneous preface paradox' discussed in n3, we are asked to imagine an author who writes a well-researched ambitious factual book in which he argues precisely that all well-researched ambitious factual books contain falsehoods (Easwaran and Fitelson, 2015, section 3). Though we are not told precisely how the author proceeds, it's natural to imagine this book as listing a series of well-researched ambitious factual books, along with the falsehoods that they have been found to contain. If that's right, then it's clear that this evidence does *not* normically support the conclusion that all well-researched, ambitious factual books contain at least one falsehood or that the author's own book contains at least one falsehood – it doesn't generate the need for a special explanation in the event that the author's book is falsehood-free. From the perspective of the normic theory, the homogeneous preface paradox presents no greater challenge than the original.

The normic theory can be easily extended to justification comparisons. Say that E normically supports proposition P more strongly than proposition Q just in case, given E, the situation in which P is false is less normal, in the sense of requiring more explanation, than the situation in which Q is false. According to the normic theory, one has more justification for believing P than Q just in case one's evidence normically supports P more strongly than Q. Given this, it is plausible that the normic theory will serve to validate Comparative Closure and the Minimum Conjunct Rule: To explain the falsity of $P_1 \wedge P_2 \wedge \dots \wedge P_n$ one must explain either the falsity of P_1 or the falsity of P_2 , ... or the falsity of P_n . Given one's evidence E, the amount of explanation required by the falsity of $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is equal to the amount of explanation required by the falsity of P_1 or the falsity of P_2 , ... or the falsity of P_n , whichever is least. According to the normic theory, the degree of justification I have for believing $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is equal to the degree of justification I have for believing the least justified of P_1, P_2, \dots, P_n . This somewhat casual demonstration can, once again, be substituted for a more formal proof, which I will detail in the next section.

Normic support is defeasible. Just because a given body of evidence provides normic support for a proposition, it doesn't automatically follow that an expanded body of evidence will do so. Suppose again that I wander into a room and notice that the wall before me appears to be red. Given that the wall appears to be red, there would have to be a special explanation in the event that the wall is not red. Suppose I then discover that the wall is illuminated by hidden red light such that it would appear to be red even if it were white. Given that the wall appears to be red and is illuminated by hidden red light, there would not need to be a special explanation in the event that the wall is not red – the new evidence, in effect, removes the need for explanation in this case. If E normically supports P, we can say that D defeats the normic support for P just in case $E \wedge D$ does not normically support P.

We can now pose the following question: Suppose a body of evidence E provides normic support for each proposition in a set $\{P_1, P_2, \dots, P_n\}$. What does $E \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n)$ normically support? That is, what is the defeating effect of learning $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$? If the normic theory is to deliver the Principle of Differential Defeat, then learning $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ must defeat the normic support for all and only those propositions in $\{P_1, P_2, \dots, P_n\}$ that were the least normically supported by E. It is far from obvious, however, that this is so. On the contrary, one might think that learning $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ should serve to defeat the normic support for *all* of the propositions in $\{P_1, P_2, \dots, P_n\}$ (giving us instead the Principle of Uniform Defeat) as it would remove the need to explain the falsity of any one of P_1, P_2, \dots, P_n . In the next section, I will argue that this hasty reasoning is mistaken. In order to do so, we will need to start thinking about normic support in a more formal way.

V NORMAL WORLDS

Suppose propositions can be ordered according to their normalcy – according to how much explanation their truth would require. Normic support can be analysed in terms of comparative normalcy relations amongst propositions: E normically supports P just in case $E \wedge P$ is more normal than $E \wedge \sim P$, and E normically supports P more strongly than Q just in case $E \wedge \sim Q$ is more normal than $E \wedge \sim P$. Given these definitions, the formal features of normic support will be determined by the formal features of the normalcy ordering of propositions.

Consider a set of propositions F, which is partially ordered by entailment, closed under disjunction and negation and which contains a 'maximal' proposition which is entailed by all propositions in the set and a 'minimal' proposition which entails all propositions in the set. The

maximal proposition can be thought of as a tautology or logical truth and the minimal proposition as a contradiction or logical falsehood. If the set of propositions F is infinite, we suppose also that it is closed under infinite disjunction. I assume that all propositions in F can be compared for their normalcy – for any two propositions, either one is more normal than the other or they are equally normal. That is, I assume that, for any two propositions in F , either the truth of one requires more explanation than the truth of the other, or their truth requires the same amount of explanation.

The maximal proposition should count as maximally normal and the minimal proposition as maximally abnormal. The truth of a tautology never requires explanation, and nothing could require more explanation than the truth of a contradiction. A disjunction will be as normal as its most normal disjunct. The only way in which $P \vee Q$ can be true is if either P is true or Q is true. To explain the truth of $P \vee Q$ is to explain either the truth of P or the truth of Q , and the amount of explanation demanded by $P \vee Q$ will be equal to the amount demanded by P or by Q , whichever is less.

If we are considering an infinite F , I also assume that there are no infinite ascending chains of increasingly normal propositions. In this case, any set of propositions will be guaranteed to have maximally normal members and we can extend the above principle to infinite disjunctions – the disjunction of any (potentially infinite) set of propositions is as normal as its most normal members. With this assumption in place it will be possible to *rank* propositions according to how normal they are: The rank 1 propositions will be the most normal ones in F . Once these are removed, the rank 2 propositions will be the most normal ones amongst those that remain, and so on. The least normal propositions in F – those that are just as abnormal as the minimal proposition – might be assigned an infinite rank¹⁰.

Given a few further assumptions, a proposition P can be modelled as a set of possible worlds – namely, the set of possible worlds at which P is true. In this case, F will be modelled by the subsets of a set of possible worlds W , with W itself serving as the maximal proposition and the empty set \emptyset serving as the minimal proposition¹¹. Disjunction will be modelled as set theoretic union and negation will be modelled as complementation in W . Given a normalcy ranking of propositions, we can derive a normalcy ranking of worlds: Let the normalcy rank of a world w be equal to the normalcy rank of the least normal proposition of which it is a member. That is, let the normalcy rank of w be equal to the normalcy rank of the least normal proposition that is true at w . The rank 1 worlds will be those that are members of only rank 1 propositions. The rank 2 worlds will be those that are members of only rank 1 and rank 2 propositions, and so on. Given this definition, as can be checked, the rank of a

¹⁰ If F is infinite, there may be propositions which are not the maximally normal members of any set reached after a finite number of steps. In this case, the most normal members of F , once all propositions with finite ranks have been removed, might be assigned, as their rank, the first transfinite ordinal ω . At this point, the process can begin anew; once these propositions are removed, the most normal amongst the remainder can be assigned rank $\omega + 1$ and so on. There may still be propositions that are not the maximally normal members of any set reached by a finite number of such further steps, in which case we can assign propositions to the rank $\omega.2$. In this way, we can reach further limit ordinal ranks – $\omega.3 \dots \omega^2 \dots \omega^3 \dots$ and we could consider uncountable ordinal ranks as well, if F is taken to be uncountably large. Though the formalism allows for transfinite normalcy ranks, it is unclear whether they can be made intelligible by the informal interpretation of normalcy in terms of the need for explanation. In any event, in the main text I restrict attention to cases in which all propositions can be assigned either a finite normalcy rank, or the single infinite rank representing maximal abnormality.

¹¹ If F is finite, it is guaranteed to be isomorphic to the power set of a set W . If F is infinite then, in addition to the assumptions already tabled, we must assume that it contains *atoms* – propositions which are distinct from the minimal proposition and entailed only by the minimal proposition and by themselves – and that every proposition, save the minimal proposition, is entailed by some atom.

proposition must be equal to the rank of the most normal worlds within it – the rank of the most normal worlds at which it is true¹².

E normically supports P just in case $E \wedge P$ is more normal than $E \wedge \sim P$. This is just to say that the most normal $E \wedge P$ -worlds are more normal than the most normal $E \wedge \sim P$ -worlds or, more simply, that P is true at the most normal E-worlds. Normic support can, then, be analysed in terms of variably strict quantification over possible worlds: E normically supports P just in case P is true at all of the most normal worlds at which E is true. Imagine the worlds at which E is true as points arrayed in space with proximity to a central point serving as a metaphor for normalcy. We can visualise worlds arranged in a series of concentric spheres radiating from that central point, as in figure 1 below:




Figure 1

The innermost sphere represents the E-worlds that are most normal, the next sphere incorporates the E-worlds that rank next in terms of normalcy and so on. The diagram depicts a situation in which E provides normic support for P_2 but fails to provide normic support for P_1 .

This analysis of normic support offers a new perspective on Closure and the normic theory. If E normically supports each proposition in the set $\{P_1, P_2, \dots, P_n\}$ then E normically supports $P_1 \wedge P_2 \wedge \dots \wedge P_n$. *Proof* Suppose E normically supports each proposition in the set $\{P_1, P_2, \dots, P_n\}$. In this case, P_1 is true in all of the most normal E-worlds and P_2 is true in all of the most normal E-worlds ... and P_n is true in all of the most normal E-worlds. It follows immediately that $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is true in all of the most normal E-worlds in which case E normically supports $P_1 \wedge P_2 \wedge \dots \wedge P_n$. **QED** The normic theory of justification validates Closure.

¹² Once propositions are modelled as sets of possible worlds, the assumption that there cannot be infinite ascending chains of increasingly normal propositions is equivalent to the assumption that any set of worlds must have maximally normal members. David Lewis (1973, section 1.4) famously considers and rejects a corresponding assumption for world similarity – which he terms the ‘limit assumption’ – though his reasons don’t straightforwardly carry over to the case of world normalcy. Without this assumption, neither propositions nor worlds could be assigned numerical normalcy ranks – but we would retain the capacity to make normalcy comparisons. This looser framework would in fact still suffice for the core aims of this section – namely, to establish that the normic theory of justification will deliver both Closure and the Principle of Differential Defeat – but the details are omitted here. For further discussion of these issues, see Smith (2016, chap. 8).

E normically supports P more strongly than Q just in case $E \wedge \sim Q$ is more normal than $E \wedge \sim P$. This is just to say that the most normal $E \wedge \sim Q$ -worlds are more normal than the most normal $E \wedge \sim P$ -worlds. If figure 2 represents, once again, the normalcy ranking of E-worlds, it depicts a situation in which E provides normic support for each of P_1 , P_2 and P_3 , but stronger normic support for P_3 than P_2 and stronger normic support for P_2 than P_1 . The most normal E-worlds in which P_1 is false are in the second sphere, the most normal E-worlds in which P_2 is false are in the third sphere and the most normal E-worlds in which P_3 is false are in the fourth sphere.




Figure 2

We might define the *degree* to which E normically supports a proposition P as the number of normalcy spheres of E-worlds throughout which P holds. Given this definition, E normically supports P more strongly than Q just in case it normically supports P to a higher degree than Q. In the above diagram, E normically supports P_1 to degree 1, P_2 to degree 2 and P_3 to degree 3¹³.

The degree to which E normically supports a conjunction $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is no lower than the degree to which it supports each of P_1, P_2, \dots, P_n . *Proof* Any world in which $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is false is a world at which either P_1 is false or P_2 is false ... or P_n is false. Let w be one of the most normal worlds

¹³ According to the normic theory (when formally developed in the way suggested here) degrees of justification, unlike of degrees of probability, are *discrete*. While the probability of two propositions can be brought arbitrarily close, while one remains more probable than the other, two propositions that differ in their degree of justification must be at least a whole step apart. This has some relevance for the kind of situation imagined in n8, in which one has equal justification for every proposition in a set apart from a single proposition for which one's justification is 'slightly' higher. If one learns that some proposition in the set is false then, according to the Principle of Differential Defeat, all of the propositions in the set will be defeated, aside from the one proposition that is more justified than the others. But what if even a single degree of justification always represents a substantial difference? The formal framework developed here is compatible with there being a limited number of possible degrees that justification could have. And, even if an infinite number of degrees is permitted, the framework can still be construed in such a way that every degree indicates a significant shift from its predecessor. These points may allay (though perhaps not altogether answer) our concerns over what the Principle of Differential Defeat predicts in a case like this.

at which E is true and $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is false. There must be some $P_x \in \{P_1, P_2, \dots, P_n\}$ such that P_x is false at w . It follows that the degree to which E normically supports $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is no lower than the degree to which it normically supports P_x . **QED** The normic theory of justification validates Comparative Closure. As can be easily checked, if $P_x \in \{P_1, P_2, \dots, P_n\}$ is false at some of the most normal worlds in which E is true and $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is false, then the degree to which E normically supports $P_1 \wedge P_2 \wedge \dots \wedge P_n$ will in fact be *equal* to the degree to which E normically supports P_x and, more generally, equal to the degree to which E supports the least supported members of $\{P_1, P_2, \dots, P_n\}$. As a result, the normic theory of justification validates the Minimum Conjunct Rule¹⁴.

If E normically supports P , D defeats this normic support just in case $E \wedge D$ does not normically support P . In this case, while P is true in all of the most normal E -worlds, P is false in some of the most normal $E \wedge D$ -worlds. D , in effect, forces us further from the most normal worlds in which E is true, and into a region in which the connection between E and P is disrupted. If the following diagram, once again, represents the normalcy ranking of E -worlds, it depicts a situation in which E provides normic support for P that is defeated by D . With evidence E , I have normic support for all and only those propositions that contain the red region – including P . Once I learn D , I then have normic support for all and only those propositions that contain the yellow region¹⁵.

¹⁴ The Minimum Conjunct Rule corresponds to one of the axioms for a *positive ranking function* (Spohn, 2009, section 2.1, 2012, section 5.2). A function r taking propositions in F into the set of natural numbers plus infinity will qualify as a positive ranking function just in case:

- (1) $r(W) = \infty$
- (2) $r(\emptyset) = 0$
- (3) $r(P \wedge Q) = \min\{r(P), r(Q)\}$

In case F is infinite, (3) may be strengthened to:

- (4) For any set of propositions Σ , $r(\wedge\Sigma) = \min\{r(P) \mid P \in \Sigma\}$

A positive ranking function that satisfies (4) is referred to as ‘completely minimative’. The degrees of normic support conferred upon propositions by a body of evidence will conform to these axioms (Smith, 2016, chap. 8, 2018, section 4). If, as discussed in n10, we assign propositions to transfinite ordinal normalcy ranks, degrees of normic support should also be capable of taking such values. In this case, the above axioms can stand, with ∞ interpreted as a kind of ‘absolute infinity’ greater than an ordinal, giving us something close to Spohn’s *conditional ordinal functions* (Spohn, 2012, pp72-73).

¹⁵ There is a clear connection between the effect of defeat, according to the normic theory, and the effect of adding conflicting information to a belief set, according to belief revision theory. On Grove’s possible worlds model of AGM belief revision theory, a belief set is represented by a set of possible worlds – the worlds at which all of the beliefs in the set are true – and is associated with a system of spheres which orders the remaining worlds with respect to their ‘closeness’ to that belief set. If, in figure 3, we interpret the spheres as representing closeness to the belief set B represented by the red region then, on Grove’s model, the yellow region will represent the new belief set that results from the addition of D to B (see Grove, 1988, section 2, partic. pp162). I won’t explore this connection further here.

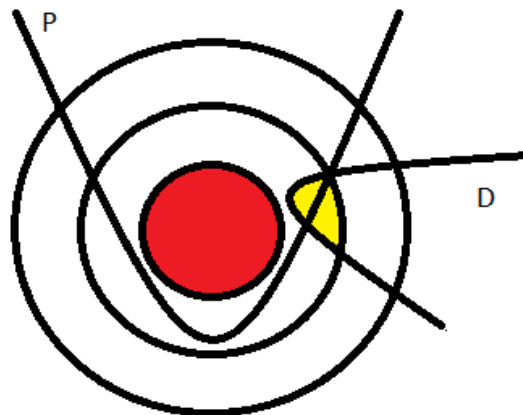


Figure 3

This diagram depicts a situation in which D defeats E's normic support for P.

Consider now the following technicolour version of figure 2:




Figure 4

Before I acquire any defeating information, I have normic support for all and only those propositions that contain the red region – a set which includes P_1 , P_2 and P_3 . If I learn $\sim P_1 \vee \sim P_2 \vee \sim P_3$, the most normal worlds that are compatible with my expanded evidence will be those in the yellow region, and I will have normic support for all and only those propositions that contain this region – a set which no longer includes P_1 , but continues to include P_2 and P_3 . If I learn $\sim P_2 \vee \sim P_3$, the most normal worlds compatible with my further expanded evidence will be those in the purple region, and I will have normic support for all and only those propositions that contain this region – a set which no longer includes P_2 , but still includes P_3 .

When I learn $\sim P_1 \vee \sim P_2 \vee \sim P_3$ this serves to defeat my normic support for the least normically supported of the three propositions – P_1 . When I learn $\sim P_2 \vee \sim P_3$ this serves to defeat my normic support for the least normically supported of the two remaining propositions – P_2 . Suppose my evidence E provides normic support for each proposition in the set $\{P_1, P_2, \dots, P_n\}$ and suppose I learn $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$. We prove in two stages the general claim that this serves to defeat the normic support for all and only those propositions in $\{P_1, P_2, \dots, P_n\}$ that are least normically supported by E . First, for any proposition $P_x \in \{P_1, P_2, \dots, P_n\}$ if there is another proposition in $\{P_1, P_2, \dots, P_n\}$ that is less normically supported by E , then $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ does not defeat the normic support for P_x . *Proof* Suppose $P_x \in \{P_1, P_2, \dots, P_n\}$ and there exists another proposition $P_y \in \{P_1, P_2, \dots, P_n\}$ that is less normically supported than P_x by evidence E . In this case, there is a world in which E is true and P_y is false which is more normal than the most normal worlds in which E is true and P_x is false. Since $P_y \in \{P_1, P_2, \dots, P_n\}$, there is a world in which E and $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ are true which is more normal than the most normal worlds in which E is true and P_x is false. Ipso facto, P_x is true in the most normal worlds in which E and $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ are true and $E \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n)$ normically supports P_x . **QED**

We can also prove the converse claim: For any proposition $P_x \in \{P_1, P_2, \dots, P_n\}$ if there is no other proposition in $\{P_1, P_2, \dots, P_n\}$ that is less normically supported by E then $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ does defeat the normic support for P_x . *Proof* Suppose $P_x \in \{P_1, P_2, \dots, P_n\}$ and there exists no other proposition in $\{P_1, P_2, \dots, P_n\}$ that is less normically supported than P_x by evidence E . In this case there are no worlds at which E and $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ are true and which are more normal than the most normal worlds at which E is true and P_x is false. Ipso facto, P_x is false at some of the most normal worlds at which E and $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ are true and $E \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n)$ does not normically support P_x . **QED**

If one's evidence E provides normic support for each proposition in the set $\{P_1, P_2, \dots, P_n\}$ then $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ defeats the normic support for all and only those propositions in $\{P_1, P_2, \dots, P_n\}$ which were the least normically supported by E . According to the normic theory of justification, if one has justification for each proposition in the set $\{P_1, P_2, \dots, P_n\}$ and one learns $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$, this new evidence serves to defeat all and only those propositions in $\{P_1, P_2, \dots, P_n\}$ that were the least justified by E . The normic theory of justification delivers the *Principle of Differential Defeat*. In the hybrid paradox, the normic theory predicts that, when I learn that the book contains an error – when I learn $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ – this will serve to defeat all and only those claims in the book that were the least justified – the least justified members of $\{P_1, P_2, \dots, P_{100}\}$ ¹⁶.

¹⁶ What if one receives only a *defeasible* report to the effect that $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ as in Ryan's third version of the preface paradox discussed in n3? Such cases are more complex, and I provide only a preliminary treatment here: Suppose one has evidence E that normically supports a proposition P and acquires evidence F that normically supports $\sim P$. Given certain conditions, this new evidence will defeat one's normic support for P just in case the degree to which F normically supports $\sim P$ is greater than or equal to the degree to which E normically supports P . The conditions in question are that E and F be equally normal propositions and that F hold in some of the most normal worlds in which E and $\sim P$ hold and E hold in some of the most normal worlds in which F and P hold. What this ensures, in a way, is that there is no extraneous normic interaction between E and F beyond that which is mandated by their levels of respective normic support for the conflicting propositions P and $\sim P$.

Ryan's third version of the preface paradox arguably fits this pattern, with one's prior evidence normically supporting $P_1 \wedge P_2 \wedge \dots \wedge P_{100}$, and the report normically supporting $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$. As mentioned in n3, Ryan suggests that, if one has evidence that justifies every claim in the book, one would not be justified in believing the report and ought to dismiss it. According to the normic theory, this diagnosis may

When accommodating new evidence, according to the normic theory, one is only required to take seriously the most *normal* ways in which the new evidence could be true, given one's prior evidence. The result proved is, in effect, a special case of this more general principle. When one learns $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ one is obliged to take seriously only those disjuncts that would be the most normal, given the existing evidence. That is just to say that one is obliged to take seriously only the falsity of those members of $\{P_1, P_2, \dots, P_n\}$ for which the existing evidence provides the least normic support.

The Principle of Differential Defeat is also an instance of another more general principle, validated by the normic theory, regarding the effect of learning $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ upon the *degree* of justification enjoyed by each proposition in $\{P_1, P_2, \dots, P_n\}$. The degree to which a proposition $P_x \in \{P_1, P_2, \dots, P_n\}$ is normically supported by evidence E is equal to the number of spheres of E -worlds throughout which P_x holds. The degree to which a proposition $P_x \in \{P_1, P_2, \dots, P_n\}$ is normically supported by $E \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n)$ is equal to the number of spheres of $E \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n)$ -worlds throughout which P_x holds. If m is the degree to which E normically supports the least normically supported propositions in $\{P_1, P_2, \dots, P_n\}$ then m is the number of spheres of E -worlds throughout which $P_1 \wedge P_2 \wedge \dots \wedge P_n$ is true, and the number of spheres that disappear when we move from evidence E to evidence $E \wedge (\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n)$. In this case, the effect of learning $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_n$ is to lower the degree of normic support for each proposition in $\{P_1, P_2, \dots, P_n\}$ by m . According to the normic theory of justification, learning that one member of a set of justified propositions is false will uniformly lower the degree of justification for each proposition in the set. In one sense, the new evidence serves to 'partially defeat' the justification for each proposition, but will only (completely) defeat – that is, lower to 0 – the justification for those propositions which were the least justified to begin with. It is only these propositions that, as it were, completely buckle under the weight of the new evidence¹⁷.

As well as validating Closure and offering viable Closure-preserving solutions to the lottery and preface paradoxes, the normic theory of justification validates the Principle of Differential Defeat and offers a viable Closure-preserving solution to the hybrid paradox. While this completes the primary aim of the paper, in the final section I turn briefly to a related topic, which also stems from the work of Marvin Backes (2019b); an objection to the normic theory of justification which focusses on the way in which it handles defeaters.

be correct provided that the degree to which the report provides justification for $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$ is lower than the degree to which one's prior evidence provides justification for the least justified of $\{P_1, P_2, \dots, P_{100}\}$. Otherwise, the report will succeed in defeating one's justification for some of the claims in the book.

¹⁷ Consider again the situation described in n8 and n13 in which one has equal justification for believing each of a series of propositions $P_1 \dots P_n$, and 'slightly' more justification for believing a further proposition P_{n+1} . If one learns $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{n+1}$ then, as discussed, the normic theory will predict that one will lose justification for believing each of $P_1 \dots P_n$, but will retain justification for believing P_{n+1} . In light of the observation made in the main text, however, the normic theory will also predict that the new information will not affect the *difference* in one's degree of justification for $P_1 \dots P_n$ and for P_{n+1} . Thus, if one's initial justification for $P_1 \dots P_n$ were just one degree lower than one's justification for P_{n+1} then, while the latter justification won't be *defeated* by the new information, it will be subject to partial defeat which reduces its degree of justification to 1. For more on partial defeat and the normic theory see Smith (2016, section 8.3)

VI THE PROBLEM OF 'EASY DEFEAT'

Suppose Helen has a peanut allergy. One day she goes into a café and orders a brownie that is labelled 'peanut free'. Helen believes that the brownie is peanut free, and this proposition is normically supported by her evidence. Suppose Helen then reads in the newspaper that a flour supplier has just announced that a bag of peanut-contaminated flour mistakenly made it into circulation. According to Backes (2019b), this will serve to defeat the normic support for Helen's belief because it provides a possible explanation as to how the brownie might contain traces of peanut, in spite of the label. If this is right, then the normic theory predicts that, once Helen reads the newspaper report, she is no longer justified in believing that the brownie is peanut free. Backes claims that this is a counterintuitive result – after all, it is extremely unlikely that any of the contaminated flour would have made its way into Helen's brownie. Backes goes on to outline several further cases of 'easy defeat', in which the normic theory allegedly makes it too easy for one's justified beliefs to be defeated. I will consider two more of his cases here:

Suppose Helen believes that she will see her friend Bob when she travels to Oxfordshire next weekend. Suppose Bob has said that he will meet her and that he is usually very reliable and trustworthy and, as a result, her evidence provides normic support for her belief. Suppose Helen then reads in the newspaper that a man in Oxfordshire has been fatally struck by lightning. According to Backes, this serves to defeat the normic support for Helen's belief, and for a similar reason to the preceding case; the new information could offer a possible explanation as to how she could fail to see Bob, in spite of her present evidence. Once again, if this is right then the normic theory predicts that, after reading the newspaper, Helen would no longer have justification for believing that she will see Bob next weekend.

Finally, suppose that Helen has, for several years, owned an apartment in New York, when she reads in the newspaper that a New York apartment was recently gutted by fire. Plausibly, Helen would still be justified in believing that she has an apartment in New York. And yet, according to Backes, the normic theory predicts otherwise – after all, this report offers a possible explanation as to how she could fail to have an apartment, despite her evidence, and thus deprives the belief of normic support.

One thing that we might observe right away is that these examples all appear to be somewhat similar to the example that drives the hybrid paradox. In the hybrid paradox, recall, I have secured evidence which normically supports each of 100 independent factual claims $P_1, P_2 \dots P_{100}$ and I then learn that one of these claims is false – $\sim P_1 \vee \sim P_2 \vee \dots \vee \sim P_{100}$. If we focus on one particular claim in the book – P_{57} say – the new evidence might be thought to offer a kind of explanation as to how this claim could be false in spite of my existing evidence. And yet, as we have seen, it is *not* inevitable that the new information will defeat my normic support for P_{57} . By the result proved in the previous section, my normic support for P_{57} will only be defeated on the assumption that it was one of the least normically supported of the 100 claims.

In Backes's examples, Helen doesn't literally learn that one amongst a set of justified beliefs is false – but she does acquire evidence which can be accommodated in a number of different ways, some of which may count as more normal, given her existing evidence, than others. Suppose one's evidence E provides normic support for P and one then learns D . If there is a proposition Q that entails D and which would be more normal, given E , than $\sim P$, then D will not defeat one's normic support for P . We might say, in this case, that Q *insulates* P from defeat by D . *Proof* Suppose E normically supports P and suppose there is a world in which E and Q are true which is more normal than the most normal worlds in which E is true and P is false. Suppose finally that Q entails D . In this case, there is a world in which E and D are true which is more normal than the most normal worlds in which E is true and P

is false. Ipso facto, P is true in the most normal worlds in which E and D are true and $E \wedge D$ normically supports P. **QED**

Consider again Backes' apartment example. On the normic theory, the information that an apartment in New York has been gutted by fire will only defeat Helen's justification for believing that she has an apartment in New York on the assumption that it would be just as normal, given her evidence, for her apartment to be gutted by fire as any other. But we have no reason to accept this – and it would take little to make it false. There are many factors that determine an apartment's vulnerability to fire. Some apartments are fitted with smoke alarms and sprinkler systems while others will lack them. Some apartments will have old, deteriorated electrical wiring, while others will have new wiring that has passed rigorous safety checks. Some apartments will have open fireplaces, while others won't, and so on. For any apartment to be gutted by fire may require explanation, but more explanation is required in the case of some apartments than others.

It would be natural to suppose that Helen is aware that her apartment has certain fire safety measures that are not present in every apartment in New York. While this evidence does leave open the possibility that Helen's apartment is the one that burned, it generates the need for further explanation in this case. Upon reading that a New York apartment has been gutted by fire, it would be natural for Helen to reassure herself with the thought that she has fire safety measures in place – a sprinkler system, wiring that has undergone safety checks etc. In this case, there is a proposition that insulates Helen's belief from defeat – the proposition that there is an apartment in New York that lacks appropriate fire safety measures and was recently gutted by fire. It would be more normal, given Helen's evidence, for this proposition to be true than for her belief to be false.

Similar remarks may apply to the Bob example (though matters are admittedly less clear-cut). Upon reading that a man in Oxfordshire has been fatally struck by lightning, it would be natural for Helen to reassure herself with the thought that Bob is relatively safety conscious and not the sort of person who would venture outside during a thunderstorm etc. The proposition that someone in Oxfordshire, more reckless than Bob, has been fatally struck by lightning would plausibly insulate Helen's belief from defeat. Once again, it would be more normal, given Helen's evidence, for this proposition to be true than for her belief to be false.

We could, of course, adjust these examples in such a way that there is *no* proposition that will insulate Helen's beliefs from defeat. We could stipulate that, given Helen's evidence, one of the most normal ways in which the fire report could be true is for her own apartment to have burned, and one of the most normal ways in which the lightning report could be true is for Bob to have been struck. Perhaps Helen knows that her apartment is highly vulnerable to fire – perhaps she has reason to believe that it is amongst the *most vulnerable* in New York. Perhaps Helen knows that Bob is well and truly reckless enough to venture out during a thunderstorm – perhaps she has reason to believe that he is amongst the *most reckless*, in this regard, in all of Oxfordshire. In this case, the normic theory would predict, as Backes claims, that the justification for Helen's beliefs is indeed defeated by the reports – but such a prediction is not obviously wrong.

With the examples fleshed out in this way, there is no available thought with which Helen could reassure herself – she would be forced to concede that her apartment is one of those that could most easily have burned, and that Bob is one of those who could most easily have been struck. Such realisations would sit very uneasily alongside the beliefs that she has an apartment in New York and that she will be meeting Bob next weekend, if she persists in holding them. Consider the tension involved in uttering 'Someone in Oxfordshire has been fatally struck by lightning and, knowing Bob, it could just as easily have been him as anyone else, but I'll be meeting Bob next weekend'. In such a

case, it's natural to think that Helen ought to fall back upon a probabilistic belief – it's *very likely* that she'll be meeting Bob next weekend or some such – until she has had an opportunity to gather more information, by contacting Bob or reading more details in the paper etc.

The problem with these two examples – a hazard of thought experiments more generally – is that the predictions and the intuitions that are supposedly being compared both turn, in part, upon details that are not explicitly supplied. In each example, Helen will have further background evidence that is potentially relevant to her belief, and to the interpretation of the new information that she receives. In each case, the kind of background evidence that we would naturally assume to be at Helen's disposal will be enough to ensure that the normic support for her belief is insulated and is not defeated by this new information.

This leaves only Backes's first example. In this case, I am inclined to think that the normic support for Helen's belief genuinely *is* defeated. When Helen reads about the bag of peanut-contaminated flour, she learns that, in all likelihood, some small proportion of baked goods in the city will contain this flour. Given her limited evidence, it would be just as normal for the brownie before her to be part of this group as any other baked good in the city. Backes claims that Helen should still be justified in believing that her brownie is peanut free – but is this really so clear? After reading the report, it would be understandable if Helen decided not to eat the brownie. And, if Helen did go ahead and eat the brownie as planned, there is some temptation to see this as a rash decision – particularly if her allergy is severe. It would also be irresponsible for Helen to assert that the brownie is peanut free or to offer some to a friend who also has a peanut allergy. Helen's new evidence does, then, make a significant difference as to how she ought to *behave* – no longer should she blithely act as though the brownie is peanut-free. A natural explanation for this is that Helen's new evidence also makes a difference as to what she should *believe* – she should no longer believe that the brownie is peanut free.

Backes describes Helen's new information as 'negligible' (pp2885) – but many, I suspect, would be inclined to take such information very seriously. If the café owners were made aware of the report, one could easily imagine them strongly advising Helen not to eat the brownie – offering her a refund or an exchange for a flour-free item. In real food contamination or tampering scares, the proportion of products affected is typically very small – but the measures taken are often drastic, including mass product recalls and dire public health warnings¹⁸.

In any case, it is not just the normic theorist who must accept that the justification for Helen's belief is defeated in this case – any defender of Closure is committed to this. Suppose that Helen were still justified in believing that her brownie is peanut free, and contains no peanut-contaminated flour. Helen should also be justified in believing that the brownie sitting unpurchased on the shelf does not contain peanut-contaminated flour, and believing the same about the blueberry muffin next to it and the cinnamon swirl at the café across the street and so on. If Helen encountered every baked good in the city, she could justifiably believe, of each one, that it does not contain peanut-contaminated flour. Given Closure, Helen would be justified in believing that none of the baked goods in the city contains

¹⁸ In September 2018, the discovery of sewing needles concealed in supermarket strawberries triggered a nationwide panic in Australia. In a reaction that is very typical of such cases, strawberries were withdrawn from supermarket shelves and consumers were strongly advised to dispose of or return any strawberries purchased. At the time of the initial reaction, only a few affected strawberries had been discovered, out of the millions consumed on a weekly basis in Australia (in the end there were over 100 reported cases). See for instance: <https://www.abc.net.au/news/2018-09-18/strawberry-fund-set-up-for-farmers-qld-premier-says/10262804> and <https://www.abc.net.au/news/rural/2018-09-19/strawberry-growers-claim-government-causing-hysteria-needle/10282070>

peanut contaminated flour. Given the newspaper report, it is clear that she does *not* have justification for believing this.

VII CONCLUSION

This paper has, in one way, been an extended exercise in exploring the consequences of preserving the Closure principle for justification. The lottery and the preface paradoxes both highlight certain commitments that any defender of Closure must be prepared to undertake – and so it is with the new hybrid paradox described by Praolini and Backes. I have argued that the primary lesson of the new paradox is that a Closure defender is under significant pressure to accept what I have called the Principle of Differential Defeat. The formal framework that I have set out demonstrates one way in which this principle could be embedded within a broad and systematic approach to epistemic justification – an approach which brings a range of further principles, such as Comparative Closure and the Minimum Conjunct Rule, as well as principles regarding partial defeat and insulation. For a Closure denier, the new paradox represents a powerful addition to one’s arsenal. For a Closure defender, the paradox is valuable in another way – for revealing more of the rich network of principles of which Closure is but one part.

Acknowledgements

This paper was presented at the Competence and Success in Epistemology and Beyond Workshop at the University of Helsinki in April 2019, the Epistemic Closure and Transmission Workshop at the University of Edinburgh in June 2019 and the Logic, Epistemology and Metaphysics Seminar at the Institute of Philosophy in February 2020. Thanks to all of those who participated on these occasions including Marvin Backes, Corine Besson, Darren Bradley, Julien Dutant, Claire Field, Giada Fratantonio, Jaakko Hirvelä, Nick Hughes, Matt Jope, Chris Kelp, Maria Lasonen-Aarnio, Lisa Miracchi, Niall Paterson, Jim Pryor, Mona Simion, Florian Steinberger, Tim Williamson and Crispin Wright. Particular thanks to Philip Ebert, Dominik Kauss and Francesco Praolini for providing detailed comments on earlier drafts of this material. Thanks also to two referees for this journal.

References

- Backes, M. (2019a) ‘A bitter pill for closure’ *Synthese* v196(9), pp3773-3787
- Backes, M. (2019b) ‘Normalcy, justification and the easy defeat problem’ *Philosophical Studies* v176(11), pp2877-2895
- Burgess, J. (1981) ‘Relevance: A fallacy?’ *Notre Dame Journal of Formal Logic* v22(2), pp97-104
- Christensen, D. (2004) *Putting Logic in its Place* (Oxford: Oxford University Press)
- Douven, I. (2002) ‘A new solution to the paradoxes of rational acceptability’ *British Journal for the Philosophy of Science* v53(3), pp391-410
- Douven, I (2008) ‘The lottery paradox and our epistemic goal’ *Pacific Philosophical Quarterly* v89(2), pp204-225

- Dutant, J. and Littlejohn, C. (2020) 'Justification, knowledge and normality' *Philosophical Studies* v177(6), pp1593-1609
- Easwaran, K. and Fitelson, B. (2015) 'Accuracy, coherence and evidence' *Oxford Studies in Epistemology* v5, pp62-96
- Eder, A. (2015) 'No match point for the permissibility account' *Erkenntnis* v80(3), pp657-673
- Evnine, S. (1999) 'Believing conjunctions' *Synthese* v118(2), pp201-227
- Grove, A. (1988) 'Two modellings for theory change' *Journal of Philosophical Logic* v17, pp157-170
- Harman, G. (1986) *Change in View* (Cambridge, MA: MIT Press)
- Kaplan, M. (2013) 'Coming to terms with our human fallibility: Christensen on the preface' *Philosophy and Phenomenological Research* v87(1), pp1-35
- Kim, B. (2015) 'This paper surely contains some errors' *Philosophical Studies* v172(4), pp1013-1029
- Kroedel, T. (2012) 'The lottery paradox, epistemic justification and permissibility' *Analysis* v72(1), pp57-60
- Kyburg, H. (1961) *Probability and the Logic of Rational Belief* (Middleton: Wesleyan University Press)
- Lewis, D. (1973) *Counterfactuals* (Oxford: Blackwell)
- Littlejohn, C. (2012) *Justification and the Truth Connection* (Cambridge: Cambridge University Press)
- Makinson, D. (1965) 'The paradox of the preface' v25(6), pp205-207
- Nelkin, D. (2000) 'The lottery paradox, knowledge and rationality' *Philosophical Review* v109(3), pp373-409
- Praolini, F. (2019) 'No justificatory closure without truth' *Australasian Journal of Philosophy* v97(4), pp715-726
- Ryan, S. (1991) 'The preface paradox' *Philosophical Studies* v64(3), pp293-307
- Ryan, S. (1996) 'The epistemic virtues of consistency' *Synthese* v109(2), pp121-141
- Smith, M. (2010) 'What else justification could be' *Noûs* v44(1), pp10-31
- Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief* (Oxford: Oxford University Press)
- Smith, M. (2018) 'The logic of epistemic justification' *Synthese* v195(9), pp3857-3875
- Smithies, D. (2012) 'The normative role of knowledge' *Noûs* v46(2), pp265-288
- Spohn, W. (2009) 'Survey of ranking theory' in Huber, F. and Schmidt-Petri, C. eds. *Degrees of Belief* (Dordrecht: Springer)
- Spohn, W. (2012) *The Laws of Belief: Ranking Theory and its Philosophical Applications* (Oxford: Oxford University Press)
- Sutton, J. (2007) *Without Justification* (Cambridge, MA: MIT Press)

Williamson, T. (forthcoming) 'Justifications, excuses and skeptical scenarios' in Dutant, J. and Dorsch, F. eds. *The New Evil Demon* (Oxford: Oxford University Press)