

# Is it ever rational to hold inconsistent beliefs?

Martin Smith

In this paper I investigate whether there are any cases in which it is rational for a person to hold inconsistent beliefs and, if there are, just what implications this might have for the theory of epistemic justification. A number of issues will crop up along the way – including the relation between justification and rationality, the nature of defeat, the possibility of epistemic dilemmas, the importance of positive epistemic duties, and the distinction between transitional and terminal attitudes.

## 1 Introduction

If we notice an inconsistency in our beliefs – a set of beliefs that couldn't all be true – then this is usually something that we would try to remedy by giving up some or all of the beliefs in question. But there are some situations in which it seems as though the rational thing to do is simply *tolerate* the inconsistency – to stick with all of the beliefs, even though we know full well that some of them must be false. Some philosophers have claimed that this is something we need to account for in our theorising about epistemic justification.

Certain theories of justification predict that the propositions one has justification for believing must always form a consistent set. One example is the *normic theory* that I have defended in previous work (Smith, 2010, 2016, 2018, 2022). According to this theory, roughly put, one has justification for believing a proposition P just in case P is true in all of the most normal possible worlds in which one's evidence holds. But if the propositions that one has justification for believing are all true at some possible world, then they must be consistent. On this theory, the justified propositions will always fit together, like pieces of a puzzle, to form a possible, partial picture of what the world is like<sup>1</sup>.

Other theories of justification don't have this feature. One prominent example is the *probabilist* or 'risk-minimisation' theory according to which one has justification for believing a proposition P just in case it is highly probable, given one's evidence, that P is true. The propositions made highly probable by one's evidence would almost always form an inconsistent set, and would not, taken together, represent any possible way that the world could be. Another theory of justification which allows for inconsistent justified propositions is a theory recently defended by Dutant and Littlejohn (2020, 2024). According to this theory, one has justification for believing a proposition P just in case it is highly probable, given one's evidence, that one is *in a position to know* P. This theory is more demanding than the probabilist theory and thus, in a sense, it doesn't allow for *quite so much* inconsistency to enter into the set of justified propositions – but it can allow for some, as we'll see. In any case, if there really are situations in which a person can rationally hold an

---

<sup>1</sup> The theories of justification endorsed by Stalnaker (2006) and by Goodman and Salow (2018, section 5) also predict that the set of justified propositions must be consistent. We could add the theories of Leitgeb (2014) and Goodman and Salow (2023, sections 4-5) too, though these theories posit an additional parameter – a question or a partition – relative to which justification is assessed. While I focus on the normic theory here, any of these other theories could, for the most part, have served in its stead. That is, the specifics of the normic theory will play no particular role in what is to come – with one possible exception which I will discuss in n8. Consistency requirements on justification have also been defended by Pollock (1983) and Ryan (1991, 1996) amongst others.

inconsistent set of beliefs, then this might be thought to favour one or other of these theories over something like the normic theory.

In this paper I will look at two alleged cases of rational inconsistent beliefs, and I will concede that one of these cases is indeed genuine. I will then argue that, once we understand how the case works, and where the rationality is *coming from*, it gives us no reason to prefer any one theory of justification over any other. On the contrary, we would expect cases of this kind to arise no matter what theory of justification we adopt. In arguing this, it will be necessary to defend a distinction between a belief being *justified* and a belief being *rational*. This might seem like a suspect move – a surreptitious attempt to change the topic perhaps. I will argue, though, that the distinction is a principled one, and one that we should all accept, regardless of our views about justification.

## 2 The cases

The most persuasive cases of rational inconsistent beliefs are, I think, to be found in the relatively recent literature – particularly Backes (2019), Praolini (2019), Littlejohn and Dutant (2020), Littlejohn (2023), Dutant and Littlejohn (2024), and Goodman (ms). But the most famous purported case of this kind is the *preface* case first described by Makinson (1965, see also Christensen, 2004, chap. 3, Foley, 2009). I will begin with a brief discussion of this case, and an explanation of why I *don't* find it particularly persuasive. Anita will be our protagonist throughout the various examples to follow... Suppose Anita has just completed a non-fiction book. She has carefully researched and checked all of the claims in the book and believes each one. But Anita is also aware that the book is long and ambitious, and that comparably long and ambitious books written by others, and perhaps even by herself, have always turned out to contain some erroneous claims. Thus, so the thought goes, she ought to believe that there are some errors in her book – and might reasonably write this in the preface. If Anita believes that there are errors in the book, and we suppose (perhaps somewhat unrealistically) that she knows exactly what the claims in the book are, then she holds inconsistent beliefs.

But is it really rational for Anita to believe that there are errors in the book? It's often a good thing to be intellectually humble, and to acknowledge that we are capable of making mistakes – but there are many ways Anita could do this without believing outright that there are errors in her book. Anita could write in the preface that there *might* be errors or even that there are *likely* to be errors – but if she goes a step further and writes that there *are* errors then, in a curious sort of way, that *wouldn't* seem all that humble. Imagine if a reviewer were to write that about Anita's book, purely on the basis of its length and ambition. That would be out of line – if the reviewer really wanted to make such a claim, they would need to put in the hard work and actually *identify* some errors<sup>2</sup>. And I don't see that the situation should be any different for Anita herself. If it's not rational for the reviewer to believe that Anita's book contains errors, it's hard to see why it would be rational for *Anita* to believe this – after all, she would have even stronger evidence for the claims in the book than the reviewer does.

Trying a different tack, if we accept the probabilist conception of justification, then that will straightforwardly predict that Anita has justification for believing that there are errors in the book, as

---

<sup>2</sup> If a reviewer did write that there were errors in the book, it would be very natural to assume that they had evidence against some particular claims – and quite surprising if no such evidence was then produced. In contrast, if the reviewer were to write 'It's very likely that there are errors in the book' then, while that might still seem rather mean-spirited, we would no longer expect any special evidence to back it up.

this proposition will be highly probable, given her evidence. And we might infer from this that it would be rational for Anita to believe it. Dutant and Littlejohn's probable knowledge theory might also make this prediction – though it depends on whether Anita's 'pessimistic inductive' evidence about other books etc. is enough for her to *know* that her book contains errors, if indeed it does<sup>3</sup>. But we can't very well appeal to these theories of justification in arguing that this is a genuine case of rational inconsistent beliefs – not if we then proceed to wield the case against rival theories. The normic theory, for what it's worth, would seem to make a very different prediction here. Any possible world in which there are errors in the book must be a world in which some particular claims in the book are false. If these claims were carefully researched and checked (which by stipulation they must have been) then this will not be a normal outcome, given Anita's evidence. On the normic theory, the proposition that there are errors in the book is not something that Anita would have justification for believing.

What we need are cases of inconsistent beliefs for which there is a strong *independent* reason for thinking that the beliefs in question are rational. I don't think that preface cases fit the bill<sup>4</sup> but, as noted above, there are some cases that plausibly do. I'm going to focus on a case described by Dutant and Littlejohn (2024, section 3). This case is similar to those described by the other authors listed at the start of this section – and might be taken as representative of the class. Suppose Anita is given a general knowledge quiz with 100 questions on a range of different topics. Being very knowledgeable about all the topics that come up, Anita answers every question with confidence and believes each answer to be correct. When she submits her quiz, she is informed by the quizmaster that exactly one of her answers is wrong, but is not told which<sup>5</sup>. What should Anita now believe?

Assuming Anita has no doubt that the quizmaster is speaking the truth, she has just three options. The first is to stop believing in all of her answers, and suspend judgment instead. Dutant and Littlejohn claim – and I tend to agree – that this would be an overreaction to the quizmaster's words, as it would involve relinquishing a lot of strongly held beliefs. If her answer to the first question is that

---

<sup>3</sup> This is one reason that I have reservations about the probable knowledge theory, and about knowledge-first approaches to justification more generally; significant questions about justification (like 'can a belief based on pessimistic inductive evidence be justified?') end up turning on corresponding questions about knowledge ('can a belief based on pessimistic inductive evidence constitute knowledge?'). But in many cases – it seems to me – the latter questions are no more tractable than the former. I will return to this issue in n6.

<sup>4</sup> Lottery cases might also be put forward as examples of rational inconsistent beliefs (Kyburg, 1961, pp197-198, Foley, 1979, 2009). Suppose a fair 100 ticket lottery has been drawn but the winning ticket is yet to be announced. Suppose Anita already believes, of every ticket, that it has lost – so she believes that ticket #1 has lost, that ticket #2 has lost and so on up to ticket #100. If Anita also believes that one of the tickets has won then her beliefs are inconsistent. But to describe this as a case of *rational* inconsistent beliefs is, I think, even more tendentious than in the preface case. Any lottery outcome is as normal as any other – which is just to say that the most normal worlds in which the lottery is run will include worlds in which ticket #1 is the winner, worlds in which ticket #2 is the winner ... and so on up to ticket #100. According to the normic theory, Anita would not be justified in believing that ticket #1 has lost or that ticket #2 has lost ... and there would be no reason to regard these beliefs as rational. When it comes to this case, the probable knowledge view will actually *join* the normic theory in its predictions. It's widely accepted that, prior to hearing the result, one cannot know that a particular ticket has lost a fair lottery (see for instance Ryan, 1996, Nelkin, 2000, Williamson, 2000, chap. 11, Ebert, Smith and Durbach, 2018, Smith, 2021, Dutant and Littlejohn, 2024). While the beliefs that ticket #1 has lost, that ticket #2 has lost ... are each likely to be *true*, they are not likely to be *knowledge*. Of the three theories that I consider here, only the probabilist theory would predict that the beliefs in this case are justified.

<sup>5</sup> These cases are often described as a variant on the original preface case, which is fair enough – but, if the coming material is on the right track, they have a very different underlying epistemic structure, in spite of their superficial similarities. An early version of a case like this is described by Ryan (1991, p304) – though she does not regard it as a case of rational inconsistent beliefs.

Sacramento is the capital of California, and her answer to the second question is that the House of Lancaster battled the House of York in the War of the Roses and so on... it doesn't seem right that Anita should just stop believing *all of these things*.

Anita's second option is to stop believing in *some* of her answers while continuing to believe in others. But which ones? If Anita is trying to determine, of a given answer, whether it could be the one mistake, there are several things that she would need to consider. She would obviously need to assess the evidence in favour of that answer and compare it with the evidence for the other answers. But there is also the matter of whether the answers stand in any *mutual reinforcement* relations – either because their contents 'hang together' or because they come from the same source etc. Even if Anita had relatively weak evidence for, say, her answers to questions 6 and 23 she might nevertheless have strong evidence for thinking that *both* of these answers would be mistaken if either one was, in which case neither answer would be a good candidate for being the *one* mistake (see Goodman, ms, section 3).

We could perhaps add the (somewhat artificial) stipulation that Anita's beliefs in her answers are epistemically independent and have exactly the same level of evidential support. In this case there would be no basis on which Anita could decide which beliefs to give up, and the decision would, from an epistemic point of view, be arbitrary. In outlining this case Dutant and Littlejohn do describe the answer beliefs as 'similarly supported', but don't commit to the stronger stipulation of epistemic independence and equal support (though a stipulation of this kind may be intended in the related cases given in Littlejohn and Dutant, 2020, section 2). But even if there are epistemic differences between the answers, and there is some principled way for Anita to make this decision, given the sheer number of factors involved, it's not something that we could reasonably expect her to do – and certainly not on the spot. So the upshot is much the same – any snap decision about which beliefs to give up would have to be an arbitrary one.

The final option for Anita, then, is to continue to believe in all of her answers, while also believing that one of the answers is wrong. This would give her an inconsistent set of beliefs. Of the three options, it's plausible to think that this is the best – it's what we could most easily imagine *ourselves* doing if placed in Anita's position. Surely it's better to tolerate the inconsistency than to abandon a whole host of beliefs or to make an arbitrary decision about what to believe. At the very least, this option doesn't seem obviously *worse* than the other two, in which case it could be a rational choice for Anita. In the quiz case, then, we have a strong reason to think that it would be rational for Anita to hold inconsistent beliefs – and it's a reason that does not presuppose any particular theory of justification. Rather, it's something that any theory of justification must somehow accommodate. For a proponent of the normic theory, that's going to be a challenge.

### **3 The principle of differential defeat**

In the quiz case, Anita is confronted with an inconsistent set of 101 propositions – that answer #1 is right (Sacramento is the capital of California), that answer #2 is right (it was Lancaster vs York in the War of the Roses)... and that exactly one of the answers between #1 and #100 is wrong. When it comes to this case we, as theorists, face a quandary that is, in a way, parallel to that faced by Anita herself. Anita has to decide which of these propositions to continue believing – and we have to decide which of these propositions Anita has justification for believing. The trouble with a view like the normic theory is that it forces us to solve *our* problem in a different way from how Anita would plausibly solve hers. As we've seen, it looks like the best thing for Anita to do is continue believing all

101 propositions. But a normic theorist can't allow that all 101 propositions are justified, because the set of propositions that one has justification for believing must always be consistent. What is causing the trouble is the following, which we might call the *Principle of Consistency*:

If one has justification for believing each proposition in the set  $\{P_1, P_2, P_3, \dots, P_n\}$  then  $\{P_1, P_2, P_3, \dots, P_n\}$  is consistent.

Other theories of justification, like the probabilist and probable knowledge views, are not committed to Consistency and would seem to have a much easier time when it comes to this case. When Anita learns that exactly one of her answers is wrong this will have little effect on the probability of each individual answer – if anything this probability might even *increase* (when Anita learns that she got exactly one question wrong she also learns that she got exactly 99 right!) As a result, the probabilist theory predicts that Anita still has justification for believing in each of her answers – and may have even more justification than she did initially. According to Dutant and Littlejohn, the probable knowledge theory will offer the same prediction, and for much the same reason; when Anita learns that exactly one of her answers is wrong this will have little effect on the probability of each individual answer amounting to knowledge – if anything this probability might increase. I think it is less clear whether the probable knowledge theory really gives this prediction, as it hinges on some non-obvious assumptions about knowledge<sup>6</sup> – but I will put this to one side here. What *is* clear, though, is that a defender of the normic theory – and anyone who is committed to Consistency – cannot say that Anita has justification for believing in each of her answers once she learns that exactly one answer is wrong. So what *should* they say?

In Smith (2022) I argued that anyone who accepts Consistency should endorse what I called the *Principle of Differential Defeat*:

If one has justification for believing each proposition in the set  $\{P_1, P_2, P_3, \dots, P_n\}$  and one learns  $\sim P_1 \vee \sim P_2 \vee \sim P_3 \vee \dots \vee \sim P_n$  (and nothing else) then, out of the set  $\{P_1, P_2, P_3, \dots, P_n\}$ , one loses justification for all and only those propositions that were the least justified prior to the discovery.

Let  $A_1, A_2, A_3, \dots, A_{100}$  be Anita's 100 answers. If Anita were told by the quizmaster that she got *at least* one answer wrong, then this would amount to learning  $\sim A_1 \vee \sim A_2 \vee \sim A_3 \vee \dots \vee \sim A_{100}$  in which case the

---

<sup>6</sup> Consider a *sensitivity* condition on knowledge of the kind defended by Nozick (1981, chap. 3). If Anita continues to believe in all of her answers, even when she learns that she got one answer wrong, her beliefs will be insensitive in the following sense: for any answer  $A_x$ , if  $A_x$  were the one wrong answer then Anita would still believe  $A_x$ . If insensitivity of this kind prevents a belief from qualifying as knowledge then Anita would not know any of the answers and it would not be probable, given her evidence, that she knows them – contrary to what Dutant and Littlejohn claim. A defender of the probable knowledge view who also accepted a condition like this would deny that the present case involves any failure of Consistency (and may even find themselves maintaining Consistency across the board). Dutant and Littlejohn are free of course to dispute this sensitivity condition on knowledge (and I would dispute it myself) – but the point is that there is room for reasonable disagreement over the knowledge-status of Anita's answer beliefs and, as a result, room for reasonable disagreement over what the probable knowledge view predicts.

Looking at this in a different way; if someone judges that Anita's beliefs in her answers, when true, could still constitute knowledge, even when she has heard from the quizmaster, then it's unsurprising that one would also judge that Anita's beliefs in her answers could still be *justified*, even once she has heard from the quizmaster. Those judgments go together very naturally. But what we want from a *theory* of justification is some insight into whether this pair of judgments is *correct*. The probabilist and normic theories will give definitive answers to this (competing ones as it turns out). But on the probable knowledge view, it seems that this is almost reduced to the status of a brute fact – something that just gets settled by fiat.

Principle of Differential Defeat would kick in and predict that she loses justification for believing all and only those members of  $\{A_1, A_2, A_3, \dots, A_{100}\}$  that were the least justified, prior to acquiring the new information. But the quizmaster doesn't (just) tell Anita that she got at least one answer wrong – he tells her that she got *exactly* one answer wrong. As a result, this case would seem to fall outside of the scope of Differential Defeat, leaving the defender of Consistency in need of some other way of determining which beliefs have their justification defeated.

But Differential Defeat turns out to be a much more flexible principle than it seems at first, and will in fact supply a verdict in this case, provided we combine it with the following *Principle of Single Premise Closure*:

If one has justification for believing P and P entails Q then one has justification for believing Q.

If Anita has justification for believing  $A_1$  then, by Single Premise Closure, she also has justification for believing that  $A_1$  is not the only false answer that she gave – justification for believing  $\sim(\sim A_1 \wedge A_2 \wedge A_3 \wedge \dots \wedge A_{100})$ . And so it is for every one of her answers – Anita has justification for believing  $\sim(A_1 \wedge \sim A_2 \wedge A_3 \wedge \dots \wedge A_{100})$  and  $\sim(A_1 \wedge A_2 \wedge \sim A_3 \wedge \dots \wedge A_{100})$  and so on right up to  $\sim(A_1 \wedge A_2 \wedge A_3 \wedge \dots \wedge \sim A_{100})$ . When Anita is told that she got exactly one answer wrong, what she learns, in effect, is that one of her answers *is* the only answer that she got wrong. That is, she learns  $(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}) \vee (A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}) \vee \dots \vee (A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})$ . No doubt that is a rather cumbersome way of putting the quizmaster's words into logical notation – but it does perfectly capture their content.

As a result, Anita's epistemic situation does, after all, fit the right template for Differential Defeat. She has justification for believing each proposition in the set  $\{\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}), \sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}), \dots, \sim(A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})\}$  and then learns  $(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}) \vee (A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}) \vee \dots \vee (A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})$  in which case, by Differential Defeat, she loses justification for all and only those members of the set that were the least justified. But if, say, Anita loses justification for believing  $\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100})$  then, by Single Premise Closure, she must *also* lose justification for believing  $A_1$ . And if she loses justification for believing  $\sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100})$  then, by Single Premise Closure, she must also lose justification for believing  $A_2$ , and so on. So the propositions in  $\{A_1, A_2, \dots, A_{100}\}$  for which Anita loses justification, when she hears from the quizmaster, are not those which are the least justified but, rather, those which *correspond* to the least justified members of  $\{\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}), \sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}), \dots, \sim(A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})\}$ <sup>7</sup>.

---

<sup>7</sup> The fact that *all* the members of  $\{A_1, A_2, \dots, A_{100}\}$  which correspond to the defeated members of  $\{\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}), \sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}), \dots, \sim(A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})\}$  are defeated will follow immediately from Single Premise Closure. To derive the result that these are the *only* members of  $\{A_1, A_2, \dots, A_{100}\}$  which are defeated we need a slightly stronger principle:

If one has justification for believing P and P, together with one's evidence, entails Q then one has justification for believing Q.

Suppose  $\sim(A_1 \wedge \dots \wedge \sim A_i \wedge \dots \wedge A_{100})$  is not one of the defeated members of  $\{\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}), \sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}), \dots, \sim(A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})\}$ . Since  $((\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100}) \vee (A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100}) \vee \dots \vee (A_1 \wedge A_2 \wedge \dots \wedge \sim A_{100})) \wedge \sim(A_1 \wedge \dots \wedge \sim A_i \wedge \dots \wedge A_{100})$  entails  $A_i$  it follows, from the above principle, that  $A_i$  cannot be one of the defeated members of  $\{A_1, A_2, \dots, A_{100}\}$ .

Consider now the following generalisation of Single Premise Closure, which we might call the *Principle of Comparative Single Premise Closure*:

If P entails Q then the strength of one's justification for believing Q is no lower than the strength of one's justification for believing P.

According to Comparative Single Premise Closure, Anita's justification for believing  $\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100})$  must be at least as strong as her justification for believing  $A_1$ , and her justification for believing  $\sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100})$  must be at least as strong as her justification for believing  $A_2$  etc. But Anita's justification for believing these corresponding propositions may also be considerably *stronger* than her justification for believing  $A_1$  and  $A_2$  etc. depending on the epistemic connections linking the members of  $\{A_1, A_2, A_3, \dots, A_{100}\}$ .

Suppose  $A_6$  and  $A_{23}$  are two of the least justified propositions in  $\{A_1, A_2, A_3, \dots, A_{100}\}$  but, along the lines suggested above, Anita has strong justification for believing that  $A_6$  and  $A_{23}$  must either both be true or both be false – that is, for believing  $(A_6 \wedge A_{23}) \vee (\sim A_6 \wedge \sim A_{23})$ . Although Anita's justification for believing  $A_6$  and for believing  $A_{23}$  is relatively weak, the corresponding propositions  $\sim(A_1 \wedge \dots \wedge \sim A_6 \wedge \dots \wedge A_{100})$  and  $\sim(A_1 \wedge \dots \wedge \sim A_{23} \wedge \dots \wedge A_{100})$  must, by Comparative Single Premise Closure, be at least as strongly justified as  $(A_6 \wedge A_{23}) \vee (\sim A_6 \wedge \sim A_{23})$  – as they are both entailed by it. Thus,  $A_6$  and  $A_{23}$  will be insulated from defeat when Anita learns that she got exactly one answer wrong.

If we stipulated that Anita's beliefs in her answers were all epistemically independent, then this would rule out connections of this kind. If we stipulated that these beliefs were also equally well justified, then Anita would have equal justification for believing, of each answer, that it is not the sole mistake – equal justification for believing  $\sim(\sim A_1 \wedge A_2 \wedge \dots \wedge A_{100})$  and for believing  $\sim(A_1 \wedge \sim A_2 \wedge \dots \wedge A_{100})$  etc. In this case, when Anita learns that there *is* a sole mistake, by Differential Defeat and Single Premise Closure, she will lose justification for believing in all of her answers. But, without these stipulations, the verdict will differ depending on the precise details of how Anita is epistemically positioned with respect to her answers, before hearing from the quizmaster.

If we accept Consistency, and accept the principles proposed in this section – all of which are validated by the normic theory<sup>8</sup> – then, given a full picture of Anita's justificatory status prior to learning that she got exactly one answer wrong, we can give a definitive solution as to which of her answers Anita still has justification for believing. But whatever the solution is, we still have the issue that it's not going to match Anita's solution to her *own* quandary. Anita's best option, as we've seen, is to continue believing in all of her answers – and this is the one solution that we *can't get*. But how surprising is this mismatch? There is at least one significant difference between our predicament and Anita's. We, as theorists, are free to *stipulate* all of the crucial facts about Anita's justificatory status,

---

<sup>8</sup> It's obvious that the normic theory will validate Single Premise Closure (both the basic version of the principle and the modified version from n7). To derive the remaining principles we will need to take the normic theory further than the brief description given in the introduction. In particular, we will need to extend the theory to cover the notion of *comparative* justification, which features in both Differential Defeat and Comparative Single Premise Closure. This can be done in a relatively natural way: one has stronger justification for believing a proposition P than a proposition Q just in case, out of all the worlds in which one's evidence holds, the most normal worlds in which P is false are less normal than the most normal worlds in which Q is false (Smith, 2010, section 2, 2016, section 2.4, chap. 5., 2022, section 4). It's obvious that the theory, so extended, will validate Comparative Single Premise Closure. To see that the theory validates Differential Defeat takes more work – the details are given in Smith (2022, section 5). The other Consistency-affirming theories mentioned in n1 could also perhaps avail themselves of these principles, though they too would need to be extended to accommodate comparative justification. I won't pursue this further here.

but if Anita herself wanted to make use of these facts to inform her decision, she would need to *figure them out* – and that, as we’ve seen, may be far from straightforward.

#### 4 Opaque defeat

Put aside, for a moment, cases of rational inconsistent beliefs and consider the following – a variant on the well-known ‘two door’ puzzle. Suppose there are two men in a room and it is part of Anita’s evidence that one always tells the truth and one always lies. Suppose Anita believes, based on testimony from a friend, that the man on the right is the truth-teller. Since her friend is generally reliable, and she has no other relevant evidence, it’s plausible that she has justification for believing this. Suppose she then overhears someone asking the man on the right the key question ‘If I were to ask the other man whether he is the truth-teller what would he say?’ to which he responds ‘No’. At this point, Anita’s justification for believing that the man on the right is the truth-teller is clearly *defeated* – for only the liar could give this answer. At this point all of the conditions for justification that have been considered in this paper will immediately cease to be met. Given Anita’s evidence, it will no longer be probable that the man on the right is the truth-teller, or that Anita is in a position to know that the man on the right is the truth-teller. And the proposition that the man on the right is the truth-teller will no longer be true in the most normal worlds in which Anita’s evidence holds.

If Anita no longer has justification for believing that the man on the right is the truth-teller, then presumably she is rationally obliged to give the belief up. That’s how defeat works. But here is a question that is not often asked in cases of defeat; *when* is she rationally obliged to give it up? Straight away? Even if Anita is an expert logician, it’s going to take her some time to process the information that she’s received and deduce that the man on the right must be the liar. For all she can tell *immediately*, the information might just as well confirm her existing belief or be completely neutral on the matter. If she’s put on the spot and asked to straightaway identify the truth-teller she should go with the testimony she has and point to the man on the right. So we can’t expect Anita to give up her belief the moment she receives the new information. But this is the moment that Anita’s belief stops being justified – for it’s the moment at which her evidence turns against it.

Even if one has justification for believing a proposition P, if one comes to believe P in a way that is not properly based on that justification, one will still be epistemically criticisable. This point is familiar – but I think we can make a symmetric point about the *cessation* of belief; even if one acquires a defeater for one’s justified belief in P, if one ceases to believe P in a way that is not properly based on that defeater, then one can be epistemically criticised for this. In order for Anita to give up her belief in a way that is properly based upon the defeater, she would need to deduce that only the liar could have given the answer she heard. And, as noted, this is going to take time. If Anita did give up her belief immediately upon receiving the defeating information then this could, at best, be the result of a fortunate guess or hunch. It would seem epistemically better for Anita to retain the belief – at least temporarily – rather than hastily abandon it for bad reasons.

This is what we might call a case of *opaque* defeat – a case in which the significance of a piece of defeating evidence is not immediately obvious to the person who receives it. The liar/truth-teller example provides a vivid illustration of the phenomenon, but it’s arguable that almost any defeating evidence will be opaque to some extent – in that it will take *some time* for a believer to process. In any case of opaque defeat there will be a lag between the loss of justification and the loss of rationality. That is, there will be an interval during which one’s belief is no longer justified but

continues to be rational<sup>9</sup>. How long this interval lasts would seem to depend on a host of factors – the content of the defeating information, the way that it's presented, one's cognitive capacities, and even practical considerations like whether one has other, more pressing, things to attend to. Life doesn't stop for one to process defeating information.

Cases of opaque defeat motivate a distinction between justified and rational beliefs. As I've noted, defenders of the normic, probabilist and probable knowledge theories all have reason to accept the distinction – and it's not just them. Let 'evidentialism' be the schematic view on which the propositions that one has justification for believing are the propositions that 'fit' one's evidence (Conee and Feldman, 1985, McCain, 2015, chap. 1, Littlejohn and Dutant, 2024, section 1). The normic, probabilist and probable knowledge theories, and many other theories of justification besides, can be portrayed as ways of spelling out the notion of evidential 'fit'. Even *encroachment* views – on which practical and/or moral stakes can make a difference to justification – are often presented in this general form, with the stakes exerting an influence on how easy or difficult it is for a proposition to count as fitting with the evidence (see for instance Pace, 2011, Basu and Schroeder, 2019, section 4.1). Once Anita hears the reply from the man on the right, the proposition that he is the truth-teller can't be described as 'fitting' her evidence on any adequate account of evidential fit – but her belief continues to be rational. And when her belief does lose its rationality – as would I think eventually happen – this need not be accompanied by any further change in her evidence or what fits with it.

Furthermore, the claim that Anita loses justification for believing that the man on the right is the truth-teller, as soon as she hears his reply, can be derived from the following very weak principle regarding the relation between justification and evidence: if a proposition P is entailed by one's evidence, then one does not have justification for believing  $\sim$ P. Finally, although we wouldn't want to insist that Anita give up her belief the moment she hears from the man on the right, it's clear that *something* changes at this point. From this point on, her belief is on 'borrowed time'. From this point on, she is subject to a rational obligation to give the belief up – even if it's unclear exactly how *time* figures in the content of that obligation. This precarious status is, I think, best explained by the hypothesis that the belief is no longer justified but continues to be rational.

Anita might be described as confronting an epistemic *dilemma*, in that she is trapped by conflicting epistemic duties (for discussion of epistemic dilemmas see for instance Hughes, 2019, 2022). To be clear, I don't think that Anita's predicament is a dilemma in the sense that she has *no rational option* – as I said, I think it is rational for her to retain her belief – but it is a dilemma in the sense that no matter what she does, she will do *some* epistemic wrong; either she will believe a proposition for which she lacks justification or she will give up a belief for bad reasons.

Before returning to cases of rational inconsistent beliefs, it's worth addressing a potential question about terminology. When it comes to the epistemic evaluation of beliefs, some epistemologists regard the terms 'justified' and rational' as being roughly synonymous. In fact, Dutant and Littlejohn treat these terms interchangeably in their defence of the probable knowledge theory<sup>10</sup>

---

<sup>9</sup> See the closely related discussion of 'rational delay' in Podgorski (2017). Podgorski observes that we cannot expect an agent to instantaneously update their beliefs upon receiving new evidence, as processing evidence takes time. Podgorski also points out that, even if an agent did instantaneously update their beliefs in precisely the way required by a piece of new evidence, there would be something rationally *undesirable* about this – if the update is instantaneous, then it could not be a *response* to the evidence (Podgorski, 2017, section 1.4).

<sup>10</sup> The paper in which Dutant and Littlejohn set out their most detailed presentation of the probable knowledge theory is *entitled* 'What is rational belief?' (Dutant and Littlejohn, 2024). In their 'Knowledge and prizes' (Littlejohn and Dutant, 2024, p288) they say '...while we understand why some authors would want to deny that rationality and justification really amount to the same thing, we do not. Those who do distinguish them should

and might insist that, given their use of the terms, Anita's belief will cease to be justified *and* rational when she hears the reply from the man on the right. On one level, the most important thing here is the distinction itself – the distinction that we find in cases of opaque defeat – and not the terms that we use to mark it out. Having said that, though, there are good reasons to use the term 'rational' to describe the status that Anita's belief continues to enjoy once its justification is defeated. The belief is rational in the sense that retaining it is the *best* option open to Anita at that point – better than giving it up on a gut feeling. If this is conceded then perhaps it doesn't matter much whether we describe this as 'rational' or choose some other term, but this description would be in keeping with the way that 'rational' is applied to choices in the practical domain – to mark those choices that are better, or at least no worse, than the available alternatives. For a belief state to be *justified* it must have a certain positive epistemic status, but a *rational* belief state, like a rational choice more generally, may simply be the best of a bad lot<sup>11</sup>.

When thinking about Anita's situation, one might instinctively reach for the familiar distinction between *bounded* and *ideal* rationality. One might say that, when Anita hears the reply from the man on the right, her belief continues to be boundedly rational – that is, rational in light of her cognitive limitations – but would not be ideally rational – would not be a belief that a suitably idealised agent would hold. Perhaps we could say that – but I have some reservations about appealing to this distinction in the present case (as useful as it may be in other domains). After all, the question that primarily interests us is what it would be rational for Anita to do, *given the options that are available to her*. If we abstract away from all cognitive limitations, and imagine an agent who can do logic instantaneously, then we end up with a new set of options and a new decision problem. In a way, there's no answer to the question of what an ideal rational agent would do if faced with Anita's decision problem, because an ideal rational agent *couldn't* face Anita's decision problem. While I agree that there may be some change in the 'kind' of rationality that attaches to Anita's belief when she hears from the man on the right, I think it might be better understood in terms of Staffel's distinction between *transitional* and *terminal* attitudes, and the kinds of rationality that attach to each (see Staffel, 2019, 2023). I will return to this in the next section.

---

take our remarks as applying in the first instance to rationality.' For the reasons given above however I don't think that the rationality of a belief can be reduced to *any* relation between its content and a body of evidence. Some epistemologists (though not Dutant and Littlejohn) have gone so far as to suggest that 'justified' and 'rational', when applied to beliefs, may be synonymous in ordinary language (see for instance Cohen, 1984, p283, Huemer, 2001, p22). There are strong linguistic grounds for disputing this however (see Siscoe, 2021, Fassio and Logins, 2023, section 6).

<sup>11</sup> There is another possible strategy for maintaining the equivalence between justified and rational belief – namely, to insist that Anita's belief continues to be rational *and* justified when she hears the reply from the man on the right. To this I would give a symmetric reply; as long as we agree that Anita's belief loses some significant positive epistemic status, then that's the most important thing – but there are good reasons to use the term 'justified' to describe the status lost. What is nonnegotiable in this case is that, when Anita hears from the man on the right, she loses justification *for believing* that he is the truth-teller. This is what the probabilist, probable knowledge and normic theories all predict, along with any theory that conforms to the evidentialist template. As a result, if Anita's belief continues to be justified past this point then we would have to sever the standard connection between justified belief and propositional justification. But I don't think we should be willing to contort the theoretical roles of the terms 'justified belief' and/or 'rational belief' just to ensure that they remain coextensive.

## 5 Rational inconsistent beliefs

What does the liar/truth-teller case have to do with quiz case? For a defender of the probabilist or probable knowledge theories the answer is ‘not much’. But if we accept Consistency then both of these cases will involve defeat – and the connections run deeper. In the liar/truth-teller case, it is initially unclear to Anita whether her new information is inconsistent with any of her beliefs, and whether it has any defeating effect. In the quiz case, it is obvious to Anita that her new information is inconsistent with her beliefs – and obvious that, given Consistency, some defeat must have taken place. But it will be highly non-obvious *which* beliefs have had their justification defeated. Either way we can’t expect Anita to react to the defeating information *immediately*.

As discussed in section 2, at the point Anita learns that she got exactly one answer wrong, she has just three options: (i) She can stop believing in all 100 answers. (ii) She can stop believing in a random selection of answers and continue to believe the others. (iii) She can continue to believe in all of her answers, even though she knows that one of them is wrong. Options (i) and (ii) will allow Anita to maintain consistent beliefs but, as we’ve seen, they both involve significant costs. We are now in a better position to appreciate exactly what these costs are.

In the same way that one can be epistemically criticised for *forming* a belief without good reason, as discussed in the last section, one can also be epistemically criticised for *giving up* a belief without good reason. The problem with options (i) and (ii) is that Anita is overwhelmingly likely to give up beliefs that are still perfectly justified, and which Anita has every reason to continue holding. And even if, by pure chance, Anita were to give up just those beliefs which have had their justification defeated, she would not be giving them up in a way that is properly based upon the defeater and, as discussed, that is *also* something that is epistemically criticisable. While option (iii) will involve holding some beliefs that are not justified, this could still be the best option on balance – or, at least, no worse than the other two.

Epistemologists have traditionally focussed on what are sometimes called *negative* epistemic duties – that is, duties to *not* believe certain things (like propositions for which we lack justification). Of the three options confronting Anita, only (iii) will involve a breach of negative epistemic duties, and that could lead one to think that, epistemically speaking, it would have to be the worst of the three. But this thought is too quick. The recent literature has witnessed a growing interest in *positive* epistemic duties – duties *to* believe certain things (see for instance Miracchi, 2019, Gardiner, 2021, Ichikawa, 2022, Simion, 2024) – and it is these duties that will be violated on options (i) and (ii)<sup>12</sup>. Whether one judges that it is rational for Anita to opt for (iii) and retain her beliefs will depend on how one weighs up these different duties. But Consistency itself does not commit us to any particular weighting.

If Consistency is correct then anyone who holds inconsistent beliefs must hold some beliefs that are unjustified. In so far as there is an epistemic duty to avoid holding unjustified beliefs, there will also be an epistemic duty to ensure that one’s beliefs are consistent. Anyone who accepts

---

<sup>12</sup> While it is conventional to categorise epistemic duties according to whether they mandate belief or an absence of belief, I think it may be more informative, for some purposes, to categorise them according to whether they mandate *changing* one’s belief state or *maintaining* it as it is. Using this scheme, the duties that are violated on option (iii) are *change* duties – duties to *give up* certain beliefs – while the duties that are violated on options (i) and (ii) are *maintenance* duties – duties to *retain* certain beliefs. At the very least, putting things in these terms may help dispel any lingering impression that the duties threatened by (iii) should automatically take precedence over those threatened by (i) and (ii). Duties to retain beliefs are discussed by Titelbaum (2016), Schroeder (2021, chap. 8) and Woodard (2022) amongst others.

Consistency is committed to this much. But there is no reason at all to think that this duty should be *paramount* – that it should automatically override one’s other epistemic responsibilities, such as not abandoning justified beliefs and not abandoning unjustified beliefs for the wrong reasons. The quiz case is, I think, nothing more than a situation in which the duty to maintain consistent beliefs is offset by other epistemic priorities. If a person did abandon some or all of their answer beliefs, just to ensure consistency, what would this say about their intellectual character? On the one hand, it would exhibit a kind of *scrupulousness* which could, in general, be an admirable trait. But here, as elsewhere, epistemic virtue is mingled with vice, for it would also show a certain *fickleness* – a willingness to give up beliefs at the merest whiff of defeating evidence.

Consistency-affirming theories of justification, like the normic theory, and Consistency-denying theories of justification, like the probabilist and probable knowledge theories, both predict that it is rational for Anita to retain her beliefs in her answers, even when she learns that one of them is wrong. But that’s not to say that the predictions of these different theories are completely identical. Notice that, according to the normic theory, a decision to continue believing in every answer would only be rational as a kind of *stop-gap*, while one figures out which beliefs to give up.

In the last section I mentioned Staffel’s insightful distinction between *transitional* and *terminal* attitudes. While a terminal attitude is the end point of a completed process of reasoning or inquiry, a transitional attitude serves instead as a kind of placeholder while reasoning or inquiry is ongoing. When Anita hears from the quizmaster, what the normic theory would appear to predict is that her answer beliefs should switch from being terminal to being transitional. But on the probabilist and probable knowledge theories, Anita’s discovery leaves the justification for her answer beliefs intact, and may even serve to *strengthen* it. If anything, then, these beliefs should become even more settled and steadfast. While I don’t propose to go into this topic in detail here, it’s unclear whether the probabilist and probable knowledge theories are getting this right. If we try to imagine ourselves in Anita’s position, it may be natural to continue believing in each of the answers, but it would also be natural to feel some *dissatisfaction* with this stance. If Anita was perfectly content with her beliefs and felt no motivation to try and restore consistency then there *would*, I think, be something criticisable about that – as if she’s being unduly complacent<sup>13</sup>.

Another way to bring out the contrast between the normic theory and the probabilist and probable knowledge theories is by appealing again to the notion of an *epistemic dilemma*. On the

---

<sup>13</sup> But would it even be *possible* for Anita to restore consistency in an epistemically responsible way? Thinking through the two door puzzle is one thing, but comparing the strength of one’s justification for 100 different quiz answers, while surveying all of the potential epistemic connections between them... that’s quite another. Even if Anita did attempt this, it is very likely that she would encounter new, relevant evidence long before she could hope to complete the task (see the discussion in Podgorski, 2017, section 2.2). But none of this is to say that Anita is under no obligation to try and make her beliefs consistent. It’s just that the most natural way for Anita to achieve this is not by reflecting on her existing evidence, but by turning her attention *outward* – asking the quizmaster which question was wrong or double checking her answers by searching on the internet, confirming with other people etc.

More generally, if we find ourselves in a situation in which our justificatory status proves particularly opaque, and resistant to our efforts to discern it, we have two options. The first is to redouble our reflective efforts – to bring all available cognitive resources to bear on the question of what we do and don’t have justification for believing and then align our beliefs accordingly. But the second option is to get outside of our own heads and seek further evidence with the express aim of *changing* our justificatory status into something *more transparent*. There are many circumstances in which the latter will be the more efficient strategy for getting our beliefs and our justificatory status to match.

normic theory, Anita is bound to contravene some epistemic duty no matter what she decides<sup>14</sup>. Even if retaining her answer beliefs is the best thing to do on balance, we might expect the decision to carry some residual regret (until such time as she is able to restore consistency). On the probabilist and probable knowledge theories, however, there is no dilemma here and Anita's decision to retain her beliefs should be guilt-free. But when we imagine ourselves in Anita's position, doesn't this *feel* more like a dilemma than an open-and-shut case?

Before concluding, we might consider again the case in which Anita has equal justification for believing in each of her answers, and in which the answers are epistemically independent. We've seen that the normic theory appears to make a particularly extreme prediction in this case – namely that Anita would lose justification for believing in *all* of her answers when she hears from the quizmaster. This might be portrayed as a serious cost of the theory, and is something that Dutant and Littlejohn rightly draw attention to (2020, section 2). What would really be 'extreme' is if the normic theory predicted that Anita should give up all of her answer beliefs in this case – that she should decide on option (i). But we can now see that this simply doesn't follow. In fact, this change to the example wouldn't lead to any change in how it would be rational for Anita to react. If Anita's answers were equally justified and epistemically independent then this might strike us, as theorists, as a very significant feature of the case – but for *Anita* this is something that would not even be apparent to her. On the contrary, it may be almost impossible for her to discern. As a result, this is something that could have no bearing on her decision – option (i) would still involve giving up a large number of beliefs for bad reasons and would still be no better, epistemically, than option (iii)<sup>15</sup>.

Even if we just *imagine* that the normic theory is correct – or some other theory that validates the Consistency principle – it's clear that cases of rational inconsistent beliefs would still arise. They would arise because of our own limitations, and because of the way in which epistemic norms can pull us in different directions. They would arise because, when it comes to epistemic decisions (just like decisions of other kinds), sometimes the best we can hope to do is to muddle through and choose whichever option strikes us as the least bad. Dutant, Littlejohn and others are right to draw attention to cases of rational inconsistent beliefs – for these cases have a great deal to teach us. But one thing that they don't teach us is the nature of epistemic justification.

## Acknowledgements

This paper was presented at the final *Varieties of Risk* workshop at the University of Edinburgh in June 2023 and at the 'Normal Meets Formal' epistemology workshop at the University of Glasgow in August

---

<sup>14</sup> Some discussions of epistemic dilemmas have focussed on conflicts between 'externalist' duties, such as the duty to only believe truths and 'internalist' duties, such as the duty to believe things that are strongly supported by one's evidence. If those are both genuine duties, then one will face an epistemic dilemma whenever one has evidence that strongly supports a falsehood. But a person in this situation could never *know* that they were facing an epistemic dilemma – if they knew that the proposition in question was false, it would no longer be strongly supported by their evidence. Some have suggested that this might be a feature of epistemic dilemmas more generally – that they are undetectable by those who are facing them (Hughes, 2019, section 2.2). But the epistemic duties that are in conflict in Anita's case would both be regarded as broadly 'internalist' and, as far as I can tell, there is nothing at all to prevent Anita from becoming aware of her dilemma.

<sup>15</sup> What if Anita really did correctly reason her way to the conclusion that her answer beliefs are equally justified and epistemically independent? Wouldn't she then be epistemically obliged to give all of the beliefs up? She would. But, for the reasons covered in n13, Anita has no epistemic obligation to attempt this reasoning process and, even if she did attempt it, it is overwhelmingly likely that her evidential situation would completely change before she could complete it.

2023. Thanks to all of those who participated on these occasions. Thanks also to two anonymous referees for this journal. Work on this paper was supported by the Arts and Humanities Research Council (grant no. AH/T002638/1).

## References

- Backes, M. (2019) 'A bitter pill for closure' *Synthese* v196(9), pp3773-3787
- Basu, R. and Schroeder, M. (2019) 'Doxastic wronging' in Kim, B. and McGrath, M. eds. *Pragmatic Encroachment in Epistemology* (New York: Routledge)
- Christensen, D. (2004) *Putting Logic in its Place* (Oxford: Oxford University Press)
- Cohen, S. (1984) 'Justification and truth' *Philosophical Studies* v46(3), pp279-295
- Conee, E. and Feldman, R. (1985) 'Evidentialism' *Philosophical Studies* v48(1), pp15-34
- Dutant, J. and Littlejohn, C. (2024) 'What is rational belief?' *Noûs* v58(2), pp333-359
- Ebert, P., Smith, M. and Durbach, I. (2018) 'Lottery beliefs: A philosophical and experimental study' *Philosophical Psychology* v31(1), pp110-138
- Fassio, D. and Logins, A. (2023) 'Justification and gradability' *Philosophical Studies* v180(7), pp2051-2077
- Foley, R. (1979) 'Justified inconsistent beliefs' *American Philosophical Quarterly* v16(4), pp247-257
- Foley, R. (2009) 'Beliefs, degrees of belief and the Lockean thesis' in Huber, F. and Schmidt-Petri, C. eds. *Degrees of Belief* (Dordrecht: Springer)
- Gardiner, G. (2021) 'Banal scepticism and the errors of doubt: On ephecticism about rape accusations' *Midwest Studies in Philosophy* v45, pp393-421
- Goodman, J. and Salow, B. (2018) 'Taking a chance on KK' *Philosophical Studies* v175(1), pp183-196
- Goodman, J. and Salow, B. (2023) 'Epistemology normalized' *Philosophical Review* v132(1), pp89-145
- Goodman, J. (ms) 'I'm mistaken'
- Huemer, M. (2001) *Skepticism and the Veil of Perception* (Lanham, MD: Rowman and Littlefield)
- Hughes, N. (2019) 'Dilemmic epistemology' *Synthese* v196(10), pp4059-4090
- Hughes, N. (2022) 'Who's afraid of epistemic dilemmas?' in McCain, K., Stapleford, S. and Steup, M. eds. *Epistemic Dilemmas: New Arguments, New Angles* (New York: Routledge)
- Ichikawa, J. (2022) 'You ought to have known: Positive epistemic norms in a knowledge-first framework' *Synthese* v200(5), pp1-23
- Kyburg, H. (1961) *Probability and the Logic of Rational Belief* (Middleton: Wesleyan University Press)
- Leitgeb, H. (2014) *The Stability of Belief: How Rational Belief Coheres with Probability* (Oxford: Oxford University Press)
- Littlejohn, C. (2023) 'Are there counterexamples to the consistency principle?' *Episteme* v20(4), pp852-869

- Littlejohn, C. and Dutant, J. (2020) 'Justification, knowledge and normality' *Philosophical Studies* v177(6), pp1593-1609
- Littlejohn, C. and Dutant, J. (2024) 'Knowledge and prizes' in Logins, A. and Vollet, J. eds. *Putting Knowledge to Work: New Directions for Knowledge First Epistemology* (Oxford: Oxford University Press)
- Makinson, D. (1965) 'The paradox of the preface' *Analysis* v25(6), pp205-207
- McCain, K. (2015) *Evidentialism and Epistemic Justification* (New York: Routledge)
- Miracchi, L. (2019) 'When evidence isn't enough: Suspension, evidentialism and knowledge-first virtue epistemology' *Episteme* v16(4), pp413-437
- Nelkin, D. (2000) 'The lottery paradox, knowledge and rationality' *Philosophical Review* v109(3), pp373-409
- Nozick, R. (1981) *Philosophical Explanations* (Cambridge, MA: Harvard University Press)
- Pace, M. (2011) 'The epistemic value of moral considerations: Justification, moral encroachment and James' 'Will to believe'' *Noûs* 45(2), pp239-268
- Podgorski, A. (2017) 'Rational delay' *Philosophers' Imprint* v17(5), pp1-19
- Pollock, J. (1983) 'Epistemology and probability' *Synthese* v55(2), pp231-252
- 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
- Schroeder, M. (2021) *Reasons First* (Oxford: Oxford University Press)
- Simion, M. (2024) 'Resistance to evidence and the duty to believe' *Philosophy and Phenomenological Research* v108(1), pp203-216
- Siscoe, R. (2021) 'Belief, rational and justified' *Mind* v130(517), pp59-83
- 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
- Smith, M. (2021) 'Four arguments for denying that lottery beliefs are justified' in Douven, I. ed. *Lotteries, Knowledge and Rational Belief* (Cambridge: Cambridge University Press)
- Smith, M. (2022) 'The hardest paradox for closure' *Erkenntnis* v87(4), pp2003-2028
- Stalnaker, R. (2006) 'On logics of knowledge and belief' *Philosophical Studies* v128(1), pp169-199
- Staffel, J. (2019) 'Credences and suspended judgment as transitional attitudes' *Philosophical Issues* v29(1), pp281-294

Staffel, J. (2023) 'Transitional attitudes and the unmooring view of higher order evidence' *Noûs* v57(1), pp238-260

Titelbaum, M. (2016) 'Continuing on' *Canadian Journal of Philosophy* v45(5-6), pp670-691

Williamson, T. (2000) *Knowledge and Its Limits* (Oxford :Oxford University Press)

Woodard, E. (2022) 'A puzzle about fickleness' *Noûs* v56(2), pp323-342