

Are there counterexamples to the consistency principle?

Clayton Littlejohn

Dianoia Institute of Philosophy, Australian Catholic University

The African Centre for Epistemology and Philosophy of Science, University of Johannesburg

Abstract: Must rational thinkers have consistent sets of beliefs? I shall argue that it can be rational for a thinker to believe a set of propositions known to be inconsistent. If this is right, an important test for a theory of rational belief is that it allows for the right kinds of inconsistency. One problem we face in trying to resolve disagreements about putative rational requirements is that parties to the disagreement might be working with different conceptions of the relevant attitudes. My aim is modest. I hope to show that there is at least one important notion of belief such that a thinker might rationally hold a collection of beliefs (so understood) even when the thinker knows their contents entail a contradiction.

0. Introduction

Consistency principles, as they're understood here, imply that each belief in a collection can only have important epistemic properties or epistemically important properties if the set of these beliefs is logically consistent. Here are three such principles. Accuracy is an epistemically important property. Being knowledge is an important epistemic property. Being an objectively epistemically desirable belief is another important epistemic property. In each case, the relevant beliefs can only have these properties if they're true. A collection of beliefs must be consistent if each belief in that collection is true. There are at least three true consistency principles.

These aren't the interesting principles. The interesting principles concern properties that even false beliefs can have. Consider being probable. This might be an epistemically important property. Each belief in a collection might be probable even if the set of these beliefs is inconsistent.¹ What about rationality? A collection of your rational beliefs might include falsehoods. Mine almost certainly would. Would a collection of your rational beliefs be consistent?

We can make a *prima facie* plausible case for thinking that a set of rational beliefs might be inconsistent using a suitable version of a preface case (Makinson, 1965). Imagine n friends visit your home over the course of the year. Each is supposed to add an interesting new fact to your visitor's book. For each of the n entries, it seems that it might be rational to believe what your friend wrote. Imagine the n th friend's entry is that one of the previous entries is mistaken.² This, too, might be worthy of belief. It thus seems that it can be rational to believe each entry to be correct even if it's certain that this collection is inconsistent.³

To help sell this suggestion, I often add the following points. First, it would be dogmatic *not* to believe the n th entry just as it would be dogmatic not to believe the first, the second, the third, etc. Second, it would be arbitrary to believe the n th entry and believe only some of the previous entries as if that would avoid the problem.⁴ Third, it would be far too externalist to say that a rational thinker will believe all the entries except the false one. Fourth, it would be far too sceptical to say that a rational thinker will believe nothing in the book (Littlejohn & Dutant, 2020). The beliefs of a rational thinker do not reflect the sceptic's unreasonable aversion to risk, the

¹ For example, if I roll a six-sided die, the propositions that I did not roll 1, that I did not roll 2, ..., that I did not roll 6 are each probable. So, too, is the proposition that I either rolled a 1, rolled a 2, ..., or rolled a 6. Each proposition is probable, but the set is inconsistent.

² We can imagine that you and each of the authors had memorised the book's contents and so knew what was and what was not in the book.

³ For dissenting views, see Adler (2002) and Ryan (1991, 1996).

⁴ One line to consider, mentioned to me by Sam Carter, is that it is indeterminate which belief is not rational to hold and determinate that at least some beliefs would not be rational to hold. My view is that this either involves a kind of objectionable arbitrariness or externalism.

dogmatist's conviction that he does not make mistakes, or an arbitrary pattern of commitment to the propositions considered.

When I've tried to make my pitch for the rational permissibility of inconsistency to students and colleagues, they're not always moved by these points. Sometimes this scepticism is due to uncertainty about the notion of belief. Is my notion their notion? Do we even have a good grip on what our own notion is? We lack mutual knowledge of a shared understanding of what belief is supposed to be. This impedes progress. In §1, I shall try to clarify one debate we might have about the consistency principle by helping us get a fix on one important notion of belief. In §2, I shall argue against the consistency principle. In §3, I shall try to strengthen my case by addressing an interesting and important objection. In this section, I shall explain why my arguments should be seen as compelling by epistemologists who see themselves as defending non-sceptical views that respect the separateness of propositions (Berker, 2013).

1. Belief

To answer questions about the consistency principle, we need to have a sufficiently clear understanding of the attitude(s) picked out by our talk of 'belief'. What is this thing I want us to discuss using this talk of 'belief'?

If we think belief is tightly connected with high degrees of confidence, rational belief might be understood like so:

Lockean: It is rational for a thinker to believe p iff it is rational for the thinker to be confident to a sufficiently high degree that p (Foley, 2009).

What's attractive about this view? Among other things, it lets us say that there might be both beliefs and credences whilst still retaining some hope that there wouldn't have to be two systems of representation and two fundamentally different sorts of mental states crowded into our heads (Sturgeon, 2008; Weatherson, 2005). The Lockean view doesn't have to say that only credence (or belief) guides action and doesn't have to say that credence and belief compete for control. Relatedly, this view predicts something that strikes many of us as obvious--that the rationality of various responses (belief included) supervenes upon an individual's values and their rational degrees of belief (Sprenger & Hartmann, 2019). It is also supported by familiar norms given some plausible assumptions about epistemic value. Given the assumption that believing is rational when believing does better in terms of expected epistemic value than suspension or the absence of belief would and given the assumption that epistemic value should be understood in terms of accuracy, we can show that the Lockean view must be correct (Dorst, 2019; Easwaran, 2016). It is also a view many non-philosophers find quite natural.

Provided the Lockean view doesn't require the highest degree of confidence, it isn't hard to see that it predicts that there should be counterexamples to the consistency principle.⁵ We'll assume that our Lockeans think that belief is sufficiently weak so as to allow that we might believe p even if there's some q such that we're more confident of q than p .

Much of the discussion of the Lockean view focuses on the question as to whether belief might require something more than a high degree of confidence. The view faces an interesting new challenge. The Lockeans I know think we only believe p if our degree of confidence in p exceeds .5.⁶ The new worry about the Lockean view is that it might be rational to believe p even if p is improbable. Suppose you think it's 40% likely that Venus Williams wins the tournament and the next most likely winners have only a 10% chance or less of winning. If asked, 'Who do you think

⁵ Although, see Clarke (2013) for a defence of a demanding version of the Lockean view and Greco (2015) for reasons to think that the requirement that belief requires probability 1 is not the recipe for scepticism that many of us take it to be.

⁶ If belief required a degree of confidence that's .5 or lower, a thinker could rationally assign p and $\sim p$ a degree of confidence of .5 and so 'rationally' believe p whilst rationally believing $\sim p$. Most of us think that pairwise consistency is a requirement on rational belief.

will win Wimbledon?' It seems both natural to say, 'I think Venus Williams will win' and to identify that as your best guess. (Granted, if asked whether Venus Williams will win or whether someone in the remaining field will win, you would answer, 'I think the winner will come from the field', but we'll explore this further below.) Fans of weak belief say that it's rational to think that p is the true answer to some question Q only if that's our best guess. They also think that it is proper or okay to think that p is the true answer to question Q when that's our best guess. Given the plausible assumption that 'think' and 'believe' pick out the same state of mind, neither thinking nor believing requires being more confident than not as our best guess in answer to a question might be nearly certain to be false (Holguín, 2022).

If the fans of weak belief think that we can think or believe things that are nearly certain to be false, they'll think belief is weaker than the Lockean do. That might suggest that they would also be open to the idea that there can be counterexamples to the consistency principle. They are not open to this idea. Consider our two questions:

Q1: Who will win?

Q2: Will Venus win?

Think of the possibilities that are compatible with assumptions operative in two conversational settings where Q1 or Q2 are under discussion. Questions will partition these possibilities and we can think of the cells of their partitions as complete answers to these questions. When Q1 is under discussion, we can think of each cell as identifying a specific winner with one player per cell. When Q2 is under discussion, the cells are more coarsely grained and it's Venus versus the rest of the field. Given the probabilities introduced above, the best guess in response to Q1 is Venus and the best guess in response to Q2 is the field. When we vary the questions under discussion but hold the credences fixed, a guess that would have been permissible had Q1 been under discussion needn't be when Q2 is under discussion. The appearance that incompatible thoughts might be sanctioned is misleading, for there's no one setting in which the incompatible guesses would be best and so no one setting in which they're both permitted (Dorst & Mandelkern, 2023; Horowitz, 2019).

We now have our second account of belief:

Weak: It is rational for a thinker to believe p iff p is the best guess (i.e., it is the most likely complete answer) in response to the contextually salient question under discussion (Goodman & Holguín, 2023; Hawthorne et al., 2016; Holguín, 2022).

Weak doesn't predict there will be counterexamples to the consistency principle because (a) the best guess always corresponds to the most likely complete answer and (b) the fans of weak belief can say that what we can be correctly described as 'believing' or 'thinking' is always relative to some context with some QUD. If we shift or change the QUD, all bets are off because the fans of weak belief can say that what's rational to think depends upon the QUD.

We have two theories of belief. Each has its merits. It's easy to see what each of these approaches says about the consistency principle. Which of these views should we adopt? We don't have to adopt either. There is a notion of belief that is important and not properly characterised by either view. We can see this once we flesh out a functional role played by something that seems similar to belief and see that neither Lockean belief nor weak belief seems to play this role.

2. Blame is not weak

Let's talk about emotions. Consider four examples:

- (i) Being pleased that the dog is in the yard.
- (ii) Hoping that the dog is in the yard.
- (iii) Fearing that the dog is in the yard.
- (iv) Being upset that the dog is in the yard.

We have four emotional responses to the possibility that the dog is in the yard. Two of these responses might be classified as positive emotions ((i), (ii)) and two as negative ((iii), (iv)). This

difference seems to track the connection between an individual's wishes and the object. Suppose I'm about to enter the yard where I'd encounter this dog if it's been left in the yard. If that's a possibility I wish for, (i) or (ii) might be true of me. If, however, I wish not to encounter dogs, (iii) or (iv) might be true of me.

There's a second distinction to draw. The difference between, say, being pleased and displeased that p indicates a difference in what's wished for, but what accounts for the difference between being pleased and hoping? Presumably, someone who hopes that p wishes for p and the same is true of the person who is pleased that p .⁷ Using terminology that is now sub-optimal, (Gordon, 1990) describes (i) as a factive emotion and (ii) as an epistemic emotion.⁸ I don't want to get into distracting debates about whether factive emotions are factive, so I'll let readers mentally insert square quotes around 'factive' as they see fit. The (apparent) fact that we can be pleased that p only if p but hope that p even if $\sim p$ is not the one that I want to focus on. The interesting difference for our purposes is connected to Gordon's observation that the kinds of reasons or grounds we might offer to explain or justify (ii) and (iv) differ from those that we might offer to justify (i) or (iii). If asked, 'Why do you fear that the dog is in the yard?', a speaker might either offer grounds for thinking that the content is possibly true (e.g., he often forgets to bring the dog in) or for thinking that there's something undesirable about the prospect (e.g., his dog is aggressive and I need to get something from his yard). With (i) and (iii), we don't see the same diversity of grounds. If I ask you why you'd be pleased that the dog is in the yard, you might say that it's good for the dog to get some fresh air or that it's good that the dog is out of the house because she hates the vacuum. What we *don't* say is, 'Well, he often has her out this time of day' or 'He told me that she's in the yard now'. The reasons or grounds offered will all have to do with the desirability of p being true rather than evidence for thinking that p is true or grounds for thinking that it's a live possibility that it is.

Why is that? Gordon attributes this to a difference in the doxastic states of someone who hopes or is pleased (or fears or is displeased). We do not offer evidence for thinking p is true when asked to explain the factive emotions because the person who asks, 'Why are you pleased that p ?' is convinced that p and the same is true of the person who is pleased. The truth of p would be common ground for the person who asks why you feel this way about p and the person who bears this emotional relation to p . The difference between (i) and (ii) or (iii) and (iv) is, in short, a difference in belief rather than wish. Hoping is wishing without believing. Being pleased requires believing that a wish has been granted.

This suggests that there's a doxastic state that, when combined with further beliefs, wishes, and desires, disposes someone who might otherwise hope/fear to be pleased/displeased. We normally describe the state that plays this role as 'belief'. I want to focus on *this* functional role. It's too often neglected in discussions of belief's functional role. It is the one most likely to put pressure on familiar views about the connection between belief and credence. If we want people to understand what belief is, directing them to think about the connection between belief, desire, and action is probably not helpful. A person known to be averse to getting wet and carrying an umbrella *might* be carrying that because they believe that it's raining, but they might carry it because they suspect that it might rain. That one carries an umbrella and is averse to rain is evidence that someone believes, but it's weak evidence. That someone is displeased that it is raining is decisive evidence that they believe. Action requires no specific beliefs and maybe no beliefs at all once we have desires and credences. Feeling can require specific beliefs or the absence of them.

Is belief as the Lockean conceives of it or weak belief suitable for playing this role of triggering emotional responses like being pleased or displeased? Are beliefs (so understood)

⁷ For a fuller account of hope and one that clarifies the subtle relations between hope and belief, see Blumberg & Hawthorne (2022).

⁸ My impression is that most authors now use the expression 'epistemic emotions' to designate emotions (e.g., curiosity) that have to do with epistemic evaluative properties rather than these ascriptions of emotions that seem to require neither knowledge nor its apparent possession of a propositional object. See Meylan (2014) for discussion.

suitable for triggering praise or blame? Do they rationalise anger or gratitude? No, I don't think so. Let's consider the kinds of tennis examples that Goodman and Holguin use. I don't know much about tennis, but I feel generally positive about the Williams sisters and not at all positive about Djokovic. If you put before me the probabilities that say, in effect, that Venus and Djokovic are the most likely to win their respective tournaments, you can clearly get me to say things like, 'I think Venus will win' or 'I think Djokovic will win'. Weak belief surfaces. How do I feel about this, you might ask. Am I pleased that Venus will win again? Am I displeased that Djokovic will win again? No. Feeling pleased or displeased comes later when the matter is settled for me. The issue isn't that the competition hasn't ended. I could be pleased in advance that I'll beat Henry Kissinger in a foot race. The tennis tournament could have been played last week and the same issues arise. I don't pay attention to these things, so if you mention that there was a tournament, give me the probabilities, and ask me who I thought won, you'll get me to say, 'I think Venus won' and 'I think Djokovic won', but you still won't make it the case that I'm pleased in one case and displeased in the other. The issues have to do with epistemic position and only have to do with time to the extent that that's connected to position.

Blame is not weak. Praise is not weak. Being pleased is not weak. We are not pleased that p just because we wish that p and p is our best guess to the QUD, so weak belief does not seem to be well suited for playing one of the central functional roles that belief is meant to play. And that suggests that there's a doxastic state that we might have tried to pick out using talk of 'belief' (e.g., when we say that the difference between being pleased and hoping is due to a difference in belief rather than taste) that is distinct from weak belief. Is this good news for the Lockean? Hardly. Suppose we get a letter reporting something we had wished wouldn't be. It could be a death, a transgression, a win for Djokovic, etc. Upon reading the letter, we will be displeased that what we feared has come to pass. Let's suppose that this required having a degree of confidence that was beyond some threshold, t . On the Lockean view, we're in the belief state that rationalises or triggers the change from fearing to being displeased because and just because, thanks to the letter, we've crossed t . If the Lockeans were offering an account of when it's rational to be in the state that rationalises these emotional responses that are stronger than mere epistemic emotions, they'd say that the key thing is crossing t . It's not.

We know that it's not because we know about the vast literature on naked statistical evidence.⁹ The literature suggests that there's a right and wrong way of crossing that threshold when it comes to belief. This is illustrated by contrasting cases like these:

Letter: Aunt Agnes wrote to tell you that she saw Uncle Vic participate in the attack on the guard.

Prison Yard: Uncle Vic was one of N prisoners in the yard when $N-1$ of the prisoners suddenly put into action a plan to attack the guard. (One of the prisoners exercising wasn't aware and wasn't involved.)

We might be convinced by Agnes's letter and take up any number of attitudes towards Vic and his actions (e.g., blame, disappointment, displeasure). Had we known the set up outlined in Prison Yard, we might have been even more confident than we would have been if we'd just read the letter in Letter, but the standard line is that it's not reasonable to blame, to be disappointed, or to be displeased given the naked statistical evidence in Prison Yard. If the Lockeans think the weakness of weak belief is all that disqualifies weak belief from playing the functional role we're

⁹ For arguments that blame requires something besides naked statistical evidence, see Buchak (2014), Littlejohn (2020), Moss (2021), Smith (2021), and Thomson (1986). Some writers think that we should explain the inability of this evidence to rationalise responses like punishment and blame because of moral encroachment. For a discussion, see Basu (2019). While agnostic about whether there is such a thing as moral encroachment, I don't think the encroachment story can explain why naked statistical evidence is insufficient in principle unless it is coupled with a further story about how the moral encroaches upon rational credence. See Gao (2019). Even if we go this far, however, I'm sceptical that the encroachment story says anything enlightening about the preface-lottery contrast or helps us understand when it is reasonable to feel certain kinds of emotional responses. For a helpful survey of the moral encroachment literature, see Gardiner (2018).

interested in, they're mistaken. Aunt Agnes might be a reliable witness, but whatever your posterior is in Letter (assuming that it's not maximal), we can choose a suitable N so that your posterior in Prison Yard is greater.

My pitch for the importance of strong belief that is distinct from weak belief or belief on the Lockean conception is that there's a doxastic state that differs from weak belief that can trigger factive emotional responses and explains why, given this person's tastes, they don't have the weaker epistemic emotion (e.g., fear as opposed to displeasure or hope as opposed to relief). This doxastic state is not rationalised by naked statistical evidence even when that evidence rationalises a degree of confidence that's greater than the degree of confidence rationalised by the kind of evidence we have in Letter. So, the Lockeans have not properly characterised the conditions under which it's rational to be in the relevant kind of doxastic state. (We can remain agnostic as to whether there is some additional state subject to norms that the Lockeans have captured.) The state that interests us might require strong evidence to be rationally held, but it doesn't just require strong evidence. It requires something else.

I'd like to explore this question about what might be needed to rationalise emotional responses of a certain kind. To do that, however, we'll need to introduce more terminology. Recall Letter. If we agree with the standard line that we can only be displeased that p if we know that p , being displeased that Vic attacked requires knowing that he did. In the good case, the letter in Letter would put you in a position to know, so when you take the contents at face value and believe that Vic attacked, your being disappointed might be objectively fitting to the situation or objectively suitable as it is both based on knowledge and inconsistent with your appropriate wishes.

What about a bad case, a case where you don't know because the letter's contents are false? We need a way of describing what's happening in this case. On a standard line, this couldn't be a situation in which you're displeased that Vic attacked just as this couldn't be a situation in which you knew that Vic attacked. Fine. You'll feel similarly in the bad case, though. Let's say that you're disappointed as if p if you're either in the good case and disappointed that p or in the bad case and otherwise similar mentally to someone in a good case. Rational belief can trigger putatively factive emotional responses (PFERs) when the relevant beliefs are true (and the emotional response is objectively fitting or suitable) and when the relevant beliefs are mistaken (and the emotional response is only subjectively suitable or fitting to the situation). We're primarily interested in the rational requirements that apply to belief, so we'll be focusing on the requirements for rationalising PFERs rather than the requirements for objectively fitting or suitable emotional responses.

Just to foreshadow things, I think that lots of epistemologists will sense an opportunity here to pitch their preferred theory of rational belief, particularly if it's one on which the conditions under which it's rational to believe involve more or less than high probability. They'll say that rational belief (as they understand it) rationalises these PFERs. I think we should test competing theories of rational belief by testing to see whether they give us a plausible story of the states of mind that rationalise PFERs, but let's not forget our history. In prior discussions of factive emotions and the doxastic differences between such emotional responses and epistemic emotions, the standard line was that factive emotions require knowledge (Gordon, 1990; Unger, 1975). This seems to be evidenced by at least the following observations. First, this seems like an unhappy thing to say:

1. Agnes was displeased that Vic joined in the attack, but she didn't know that he was involved.

According to some, this is an unhappy thing to say because it can only be true that someone is displeased or pleased that p if they know that p . This also seems like a bad thing to say:

2. Agnes was displeased that Vic joined in the attack. Not only that, she knew that he assaulted someone.

The second bit seems redundant given the first, but that's an indication that the information contained in the second part is contained in the first. That's some evidence that the first part is true only if the second part is. Relatedly, this seems like a bad question to ask:

3. Agnes was displeased that Vic joined in the attack, but did she know that he was involved?

I think this provides us with a clue about what belief is supposed to be (i.e., which beliefs are objectively suitable and which ones are subjectively suitable). If it's supposed to be suitable for playing a functional role and that includes, *inter alia*, triggering seemingly factive emotions, we might think that beliefs do what they're supposed to do only when they're knowledge. Much in the way that there's something wrong or inappropriate with being upset when, as you'd put it, *p*, there's something wrong or inappropriate with believing that something upsetting has happened when you don't know that it has. The preferred development of a knowledge-first theory of rationality favours a particular line on consistency, but we'll hear from the other side on this debate.

3. Against Consistency

Let's consider some new letters:

Letter: Aunt Agnes wrote to tell you that Vic has passed away.

Letters such as this are potential sources of knowledge. Suppose we're in the good case. In light of the letter, we can imagine that you're quite upset that Vic has passed. Grief is a fitting response, objectively and subjectively.

If that's true in this case, it could be true in these cases, too:

Letter 2: Grandma Brenda wrote to say that Charis passed away.

Letter 3: Uncle Chuck wrote to say that Dinesh passed away.

...

Letter 100: Cousin Zelda wrote to say that Alice passed away.

With each letter, a new opportunity for knowledge and rational belief. With each new letter, we get new reasons to grieve.

We can imagine that each letter comes from a trusted source and a reliable informant. We're going to explore some questions about what attitudes might be rational to hold when we learn further things about these letters and the beliefs we have about their contents, but let's pause to think about what makes it rational to believe in each instance the contents of these letters considered on their own. Suppose we want to explain why this kind of testimony provides an opportunity for rational belief when naked statistical evidence wouldn't. One going theory is that the kind of naked statistical evidence we might get in a lottery case, say, rationalises a high degree of confidence but doesn't make outright belief rational because the situations in which that evidence supports a falsehood aren't less normal than the situations in which it supports a truth (Smith, 2016). This is the normic support view. When it comes to the letters, the rough idea is that the letters can provide normic support as evidenced by the fact that their falsity would require a special kind of explanation that isn't required when naked statistical evidence supports a falsehood. It's a helpful foil for two reasons. First, it says (correctly, I think) that naked statistical evidence does not make it rational to believe in the sense we're interested in. Second, it predicts that there is a consistency requirement on rational belief. In my view, it gets something importantly right that the weak belief view and Lockean view get wrong whilst also getting something wrong because it imposes a consistency requirement. There's a related theory says that it's evident that an explanatory connection between fact and belief is missing in the lottery cases, but not the testimony cases (Nelkin, 2000). The lack of rational support in the lottery case can be explained in terms of the obviousness of the lack of this support. I mention this view to set it aside because whilst this view is similar to the normic support view, it is difficult to see what it predicts about consistency.

An entirely different approach is a knowledge-centred approach. In the naked statistical evidence cases like Prison Yard, the evidence both warrants a high degree of confidence and makes it certain that you're not in a position to know. If we think that believing rationally requires believing in a way that's sufficiently similar in subjective respects to a possible case of knowledge (Bird, 2007; Ichikawa, 2014; Rosenkranz, 2021), it's easy enough to see why the kind of support

that testimony provides differs in kind from the support naked statistical evidence provides. Letters can transmit knowledge on all the going theories of knowledge. Naked statistical evidence cannot provide knowledge on many of the most plausible theories of knowledge.¹⁰ That's the difference. There's a version of the knowledge-centred view that I prefer, a view on which rational beliefs are rational because they're sufficiently likely to be knowledge (Littlejohn & Dutant, forthcoming). Testimony makes it rational to believe when it is rational to be sufficiently confident that the testimony is a basis for knowing. Because it is certain that naked statistical evidence is not a basis for knowing, it is not a basis for rational belief. These knowledge-centred views seem to differ from the normic support view in the following way. In a suitably described preface case, it seems that we might have an inconsistent set of propositions where each proposition is such that it is highly likely to be something we do know or could know. Thus, on our preferred knowledge-centred view (and unlike the normic support view), the lottery and the preface differ in an important way. Because we can know quite a lot in preface cases and nothing in lottery cases, we can have rational beliefs in the former and not in the latter.

Let's return to the letters. Letter rationalises grief. The same goes for Letter 2. The same goes for letters 3 through 100. Letter rationalises grief only if it makes it rational to believe that a loved one has been lost. That holds for the other letters, too. In this setting, everyone seems to agree to this much *and* agree that the reason that the letter rationalises grief is *not* that it rationalises a higher degree of confidence than one warranted in a lottery case. When we believe the contents of any letter, there's some non-zero chance that we'll believe a falsehood and/or believe something we don't know. Bearing this in mind, imagine that each author is equally reliable and equally deserving of our trust and let's suppose that your confidence in each case that the contents of the letter are accurate is .95. Given this and assuming that the authors reports were suitably independent, the expected number of errors in this collection of errors is 5. The actual number might be smaller. It might be larger. That's the expected number, however.

Here's my first observation:

Expected Error Tolerance: It can be rational to believe p even when it is part of a large set of similarly supported beliefs where the expected number of errors contained in this set is much larger than one.

Upon pain of scepticism, Expected Error Tolerance must hold. We should note, however, that tolerance has its limits:

Limited Expected Error Tolerance: It cannot be rational to believe p when it is part of a large set of similarly supported beliefs where the expected number of errors is greater than the expected number of accurate beliefs.

This is a consequence of the idea that a rational belief is, *inter alia*, more likely than not to be correct in light of the thinker's information. Combined, we get the idea that the set of things rationally believed is not the set of things that are completely certain and that information that makes you think it's less likely than not that p should lead you to abandon that belief.¹¹

What about actual error? What happens when you get evidence that establishes to your own satisfaction this?

Concession: One of the beliefs based on the letters received is mistaken.

¹⁰ Although, see Engel (2021) and Gibbons (2013) for dissenting views on knowledge and lotteries.

¹¹ In Littlejohn & Dutant (forthcoming), we argue that an advantage of the Lockean view and our preferred theory of rational belief is that it explains why rational belief is connected to expected desirability and why rational tolerance for the expectation of error has its limits. On our view, rational beliefs are rational because they're sufficiently likely to be knowledge. We think it is a desiderata on a theory of rational belief that it actually explains these tolerance principles. Of course, these principles are controversial. See Smith (forthcoming) for a defence of the view that it can be rational to believe the highly improbable.

The Lockeans will say that the evidence in light of which we know the expected number of errors exceeds one is the kind of evidence that makes it rational to believe Concession, but we're not Lockeans. We might think, for example, that naked statistical evidence wouldn't warrant Concession even when, in light of it, we acknowledge that the number of expected errors greatly exceeds one. Testimony, however, could provide rational support for Concession much in the way that it provided rational support for beliefs about your various distant relatives via the death letters. Should we say that testimony can both support believing the contents of each letter and Concession when the thinker knows that the content of some letter must be mistaken if Concession is correct?

My first point in favour of rational inconsistency is based on a thought and an observation. The thought is that clear cases of knowledge are plausible cases of rational belief. The observation is this:

... there are important differences between the lottery and the preface. An especially noteworthy one is that in the preface you can have knowledge of the propositions that make up your book whereas in the lottery you do not know of any given ticket that it will lose. This difference, however, is to be explained by the prerequisites of knowledge, not those of rational belief (Foley 2009: 44).

Like Foley, I think that the beliefs in the contents of the letters can be clear cases of knowledge even when we have testimony that supports Concession. And, like most epistemologists, I have internalist instincts about rationality. If there is one false letter in the pile, believing this letter is not less rational than believing the other letters. If we put this together, our intuitions about knowledge support the idea that there can be inconsistent sets of propositions rationally believed. Grief can be fitting even when the expected number of cases where grief backed by similar grounds is not objectively fitting is much greater than one. It can remain fitting even when the known number of cases where it is not objectively fitting is one. What goes for grief goes for belief. Grief and belief are not fitting in a lottery-type situation, but they can be subjectively and objectively fitting in preface-type cases.

The second argument for thinking that there can be inconsistent sets of propositions rationally believed concerns Concession and Expected Error Tolerance. Let's imagine two ways that the letters case could have developed:

Mere Expectation: The letters are read and committed to memory. The thinker realises that there's some small chance that each letter is mistaken. Given their information, the expected number of errors is 5.

Conviction: As before, but we add that we're told by a reliable source that one (and only one) letter contained an error.

Upon pain of scepticism, it must be possible to rationally maintain each of the beliefs in Mere Expectation. What about Conviction? Here, some authors will say that we experience a kind of defeat. When the evidence warrants believing outright that one belief is mistaken, the beliefs cannot each be rationally held. This proposal faces two objections.

Remember that Expected Error Tolerance is common ground accepted by those who defend consistency requirements and those who reject them. Mere Expectation should be seen as a case in which each of the beliefs about the contents of the letters can be rationally held. The fans of consistency have to say that rational support is defeated or subverted to such a degree that it's not possible to rationally hold these beliefs in Conviction, but the 'news' we have about our beliefs is better in Conviction than Mere Expectation. With respect to each belief about the contents of each letter, your confidence that the belief in question has the properties that make it objectively epistemically desirable should be greater in Conviction than Mere Expectation (e.g., whilst your confidence in Mere Expectation that each letter is a potential source of knowledge might be a maximum of .95, your confidence in Conviction could be .99). Think about this in terms of the

news we would rather receive about our beliefs. We would prefer to learn that we've made only one mistake in coming to believe some large number of propositions on diverse subject matters than learn that the expected number of errors in this collection greatly exceeds that. If that is right, we should accept that in the Conviction case, we can rationally retain the beliefs that were rationally formed before we learned about the mistake. But then, we should reject the consistency principle.

Here's a second line of argument for thinking that we should drop the consistency requirement. Consider three positions we might take with regard to this requirement and our tolerance principles:

	Consistency Requirement	Expected Error Tolerance	Limited Error Tolerance
My view	No	Yes	Yes
Consistency (a)	Yes	Yes	Yes
Consistency (b)	Yes	Yes	No

We're assuming that every view under consideration is an anti-sceptical view and that every anti-sceptical view accepts Expected Error Tolerance. The interesting issue for those who defend the consistency requirement on collections of rational beliefs is to choose which tolerance principles they'd want to accept and to explain these principles.

Limited Expected Error Tolerance seems like an eminently reasonable principle. It's easily vindicated by any view according to which it is only rational to believe p if it's more likely than not that p . Thus, we might want to focus on the first set of options, one that upholds the two tolerance principles as well as the consistency requirement. The oddity of this combination of commitments, to my mind, is this. If this view is correct, either a rational belief or a body of evidence that rationally supports such a belief about the presence of falsehood is a rational toxin too powerful for continued rational support. (That, in essence, is a commitment of the consistency requirement.) My sense is that epistemologists spend far too much time thinking about the defeaters that come by way of full belief or conviction and not enough by way of credences or possibilities we assign positive probabilities to. If we accept Limited Expected Error Tolerance, we have to acknowledge that the possible presence of some rational toxins can itself sufficiently noxious to be a rational toxin in its own right. What we all need is some story about when and why evidence that indicates the possible presence of some rational toxin is itself too toxic for rational belief.

Here's a natural story, to my mind. If, in light of the evidence, the probability that a belief has an objectively undesirable epistemic property (e.g., that the belief is inaccurate) is too great, it's not rational to hold that belief. We don't need evidence that warrants the outright belief that some belief has this property to make it irrational to maintain this target belief. If, however, the probability that the belief has this property is too low, the possibility that the target belief has this property is rationally tolerable. This story is incompatible with the pro-consistency view we're considering. On the pro-consistency view, it is *not* rationally acceptable to believe the contents of each of the letters when the evidence warrants *conviction* that there is one error but states that it is rationally acceptable to believe when, in light of the evidence, the expected number of errors greatly exceeds one. In Conviction, the probability that a belief has an objectively undesirable property could be lower than it is in Mere Expectation. The pro-consistency view has to reject the probabilistic story that seems quite natural for explaining why tolerance for expected error has limits.

4. An argument for the consistency requirement

Let me conclude by addressing an argument for the consistency requirement. As I see it, we should accept the following claims:

Small Risk: It can be rationally permissible to believe even if there's a small risk of violating some objective epistemic norm (e.g., the truth norm, the knowledge norm).

Accumulated Risk: If each new belief carries only a small risk of violating some objective epistemic norm, we can continue to form new beliefs even if on the aggregate the probability of some such violation gets near certainty and the expected number of such errors greatly exceeds one.

Known Violation: It can be rationally permissible to form (and retain) said beliefs even when it's known that precisely one of them violates the truth norm.¹²

Smith disagrees. He seems to accept Small Risk and Accumulated Risk and so rejects Known Violation.¹³ To explain this choice, he asks us to consider this case:

Looters: Suppose an electronics store is struck by looters during a riot. 100 people carry televisions from the store, while the transaction record at the cash register indicates that only one television was legitimately purchased, though no receipt was issued. Suppose Joe is stopped by the police while carrying a television from the store and is promptly arrested and charged with theft. When his case comes to trial, the prosecution presents the evidence that Joe is one of the 100 people who carried a television from the store during the riot and that 99 of those people are guilty of theft (2021, 93).

Smith insists that in this case, we cannot convict Joe or any of the others, arguing as follows: if it would be permissible to convict Joe, it would be permissible to convict the rest. But we *know* that one of them is innocent, so we must refrain from convicting each of them. (We seem to agree that convicting only some of the people in the store is not an acceptable response.)

It might seem to some that there are two potential problems with convicting in Looters. The first is that we'd be using naked statistical evidence to serve as the basis for conviction. The second is that if we convicted each person, we would perform a collection of actions that, *inter alia*, would involve convicting an innocent person. Smith uses this second (putatively) intolerable consequence to explain why the first is intolerable. The problem I have with his preferred explanation is that I think we often *should* perform a series of actions that, *inter alia*, means we'd know we'd convicted an innocent person. Why does Smith think we shouldn't do that?

He thinks that convicting all 100 is, "tantamount to deliberately convicting an innocent person" (2001, 95). He thinks the situation is, in some respects, similar to cases from ethics classes where some small-town sheriff contemplates knowingly framing someone to restore order and prevent a mob from killing 99. He acknowledges that the analogy isn't perfect:

There are, of course, some notable differences between the two cases. In the ethics test case, as usually described, there is a particular, known individual we can choose to convict. In the present example [of Looters], we have the choice of convicting some innocent individual or other, without knowing who the individual is. The former, we may say, is a matter of deliberately convicting an innocent *de re* while the latter is a matter of deliberately convicting an innocent *de dicto*. While this is worth pointing out, it has no obvious relevance for the

¹² For a helpful overview of options and the difficult choices we have to make when we think about the aggregation of the risk of violating a non-consequentialist norm, see Lee-Stronach (2021).

¹³ Smith seems to think that risk (understood probabilistically) does not quite have the normative significance I take it to. See Smith (forthcoming a). I worry that views that use normalcy or modal notions (e.g., (Whiting, 2022)) to rank options might clash some very attractive principles (e.g., stochastic dominance) about how we ought to choose in the face of uncertainty (Littlejohn, 2023).

permissibility of these actions. The intention element of criminal offences can be satisfied by intentions that are either *de re* or *de dicto* in the sense defined here. To satisfy the intention element of, say, murder or theft, it is sufficient that one intend to kill/take the property of some person or other (2001, 95).

Still, though, he thinks the analogy is close enough and that much as we shouldn't convict a collection if it contains someone known *de re* to be innocent, we shouldn't perform a collection of convictions if it's known *de dicto* that an innocent would be convicted.

Consider, then, a moral argument and an analogous epistemological one:

Against Convicting Given Knowledge of Innocence

M1. It would be wrong to convict given knowledge *de re* of innocence.

M2. There is no moral difference between convicting given knowledge *de re* of innocence and knowledge *de dicto* of innocence.

MC. Thus, it is wrong to convict given only knowledge *de dicto* of an innocent being convicted.

Against Believing Given Knowledge of Error

E1. It would be wrong to believe given knowledge *de re* of error.

E2. There is no epistemic difference between believing given knowledge *de re* of error and knowledge *de dicto* of error.

EC. Thus, it is wrong to believe given knowledge *de dicto* of error.

The argument for EC doesn't tell us whether it is also intolerable to hold a set of beliefs when the expected number of errors in that collection is much greater than one, but we know that if we endorse this reasoning and reject Known Violation, we'll either have to say that we should prefer belief sets with a greater number of expected errors to one that contains only one or embrace the sceptical view that says that it's not rational to hold large belief sets when the expected number of errors it contains exceeds one. Neither response strikes me as plausible, so I'll need to address the epistemic argument.

The first thing I want to note in response to Smith's argument is that his crucial assumption that there's no normative difference between knowing *de re* and *de dicto* that an innocent will be convicted is quite controversial. Seeing why that is will hopefully be instructive. Recall Scanlon's example involving the transmitter room (1998, 235). A technician is tangled in wires and will die a very painful death from electrical shocks if we do not shut down the broadcast. If we do stop the broadcast to free our technician, football fans will miss the first fifteen minutes of the World Cup final. Scanlon wants us to compare the complaints that the technician could raise if we keep the broadcast going and the complaints of the fans if we cut the power to free the technician:

- On the one hand, if we compare the complaints of the technician and of one fan (e.g., Oliver Miller from Ruislip), it's clear that the technician's complaint is weightier.
- On the other, if we combine the complaints of each fan into one aggregate complaint (where each additional complaint adds the same weight), the technician's complaint would no longer be the weightiest.

Scanlon thinks we should not 'combine' complaints and rather thinks only the complaints of individuals should be compared in trying to decide what to do. The aggregate of football fans is not, he thinks, an object of moral concern. If so and the only objects of moral concern are distinct individuals, perhaps respecting the separateness of persons requires respecting this constraint that we focus on individuals and only the interests of individuals.

Suppose we wanted to formulate a view that both respected Scanlon's individualist constraint and told us what to do in the face of uncertainty? A natural first idea is that a complaint's weight might be determined by the magnitude of harm and the probability that it's suffered by some individual. This seems to lead naturally to a view that *predicts* that the difference between

knowing *de re* that someone will suffer a harm or loss and knowing that *de dicto* has considerable moral significance. Consider one more example, modified from (Otsuka, 2015). A pathogen carrying comet will strike Nebraska and cause a plague that will kill millions if we do nothing. We can destroy the comet using a missile. We can destroy it using a laser. Either way, we prevent the pathogen from killing millions. If we fire the missile, however, a piece of the missile will hurtle towards Boca Raton and strike Bob Belichuk costing him his right arm. If we fire the laser, the dust from the comet will drift over California. This tiny bit of dust will be inhaled by one person and will be fatal to them, but it's completely random who will be the unlucky person who inhales it.

- On the one hand, if we compare the complaints of Bob (I'll lose an arm!) and Brenda Oglethorpe of Cool, CA (There's a 1/40,000,000 chance that I'll inhale this dust!), it's clear that Bob's complaint is weightier.
- On the other, if we combine the complaints of each of the 40,000,000 Californians, it's clear that Bob's complaint is less weighty.

If the only locus of moral concern is an individual and not some abstract California resident who somehow never once resided in California, we might focus on just the comparative weightiness of the complaints of Bob and each Californian like Brenda. If so, we might fire the missile instead of the laser.¹⁴

We now find that many epistemologists endorse the separateness of propositions and so presumably accept something like Scanlon's idea that the loci of epistemic concern will be individual beliefs and their properties.¹⁵ When we think about the letters case, the situation looks like this:

- If (for each proposition) we compare the 'complaints' about suspension (That's almost certain to be something I'd know!) and about belief (There's a tiny chance that's normally rationally tolerable that this belief will be false!), it's clear that the complaints about suspending are weightier.

Belief compares favourably to suspension in each case. The second part of the separateness idea is the negative claim that we cannot simply aggregate collections of beliefs and think of such collections and their properties as a second locus of epistemic concern.

If this is correct, it tells us what matters and what doesn't when we learn that a collection contains a falsehood. What's relevant is the risk in each individual case that a belief has an undesirable property. Further information about the collection (e.g., that it contains an error, that the expected number of errors is this or that) has no relevance to questions about whether to believe or suspend in any particular case. If we believe in the separateness of propositions and understand that in the way I've suggested, we should not be surprised that *is* an epistemic difference between believing given knowledge *de re* of error and knowledge *de dicto* of error. Believing given knowledge *de re* of error is believing something known to be false. Believing given knowledge *de dicto* of error (except in the degenerate case where we're dealing with one belief) does not compel us to see any 'choice' between belief and suspension as one in which we'd take on a belief as having

¹⁴ If this seems outrageous, note that Bob would gladly trade places with a Californian if the choice was the sure loss of an arm and the 1/40,000,000 chance of losing his life. And if that doesn't help, remember that I'm describing a moral view, not endorsing it.

¹⁵ Whilst I do not agree that the separateness point causes trouble for reliabilism, *per se*, I do believe that the point causes trouble for epistemic consequentialism (Littlejohn, 2012), but perhaps not for every form of teleological view (Littlejohn, 2018). The separateness of propositions is not easy to formulate and requires defence. For scepticism about the idea, see Sharadin (2018) and Talbot (2014). In my view, epistemic norms are silent on whether trade-offs are permitted because they are concerned with the properties of individual beliefs only, though I agree with Sharadin that a parallel view in ethics is impossible to defend. Even some of the most important critics of aggregation in ethics think that some trade-offs are required or permitted (Taurek, 2021).

an objectively undesirable property. This *de dicto* knowledge matters only insofar as it gives us reasons to increase or decrease our confidence that some belief or suspension has an objectively undesirable property. I would recommend comparing belief to suspension in light of this information on the basis of their respective expected objective epistemic desirability. In doing so, we wouldn't have any reason to think that accepting Small Risk and Accumulated Risk compels us to reject Known Violation. We should accept all three.

Conclusion

I've shown that there is an interesting notion of belief that differs from Lockean belief and weak belief, one that plays an important role in rationalising emotional responses. This notion is interesting because it seems that rationally being in this state requires something besides or beyond a high degree of confidence that the target proposition is true, as evidenced by the fact that naked statistical evidence might warrant a high degree of confidence that, in turn, rationalises hoping or fearing that p without being pleased or displeased by the (apparent) fact that p . Belief, so understood, might be rational but mistaken and it's an interesting question whether a collection of such states might each be rational when it's certain that they cannot all be correct. I've argued that a thinker might rationally hold a collection of such states even if they're known to be inconsistent and tried to show that this isn't surprising if we accept the separateness of propositions. If we accept that, then the knowledge that a collection contains a falsehood only has an indirect bearing on the rational status of individual beliefs. My hunch is that it bears on the rationality of full belief only by giving us evidence that the belief does (or does not) have desirable properties. In my view, such information about the collection only makes beliefs irrational if it's a sufficiently strong indicator of ignorance (i.e., only if the probability that this belief is knowledge is insufficiently high) (Littlejohn & Dutant, 2021). Since being part of a large collection of otherwise promising beliefs known to contain a single falsehood is often *excellent* evidence that a belief constitutes knowledge, it seems that such knowledge is no real threat to the rational status of these beliefs.¹⁶

References

- Adler, J. E. (2002). *Belief's Own Ethics*. MIT University Press.
- Basu, R. (2019). What We Epistemically Owe to Each Other. *Philosophical Studies*, 176(4), 915–931. <https://doi.org/10.1007/s11098-018-1219-z>
- Berker, S. (2013). Epistemic Teleology and the Separateness of Propositions. *The Philosophical Review*, 122(3), 337–393. <https://doi.org/10.1215/00318108-2087645>
- Blumberg, K., & Hawthorne, J. (2022). A New Hope. *Journal of Philosophy*, 119(1), 5–32.
- Buchak, L. (2014). Belief, Credence, and Norms. *Philosophical Studies*, 169(2), 1–27. <https://doi.org/10.1007/s11098-013-0182-y>
- Clarke, R. (2013). Belief Is Credence One (in Context). *Philosophers' Imprint*, 13, 1–18.
- Dorst, K. (2019). Lockeans Maximize Expected Accuracy. *Mind*, 128(509), 175–211. <https://doi.org/10.1093/mind/fzx028>
- Dorst, K., & Mandelkern, M. (2023). Good Guesses. *Philosophy and Phenomenological Research*, 105(3), 581–618.
- Easwaran, K. (2016). Dr. Truthlove Or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Nous*, 50(4), 816–853. <https://doi.org/10.1111/nous.12099>
- Engel, M. (2021). Lotteries, knowledge, and inconsistent belief: Why you know your ticket will lose. *Synthese*, 198(8), 7891–7921. <https://doi.org/10.1007/s11229-020-02555-w>

¹⁶ I presented an early version of this paper to the 18th annual Episteme conference in Zanzibar. I want to thank Jennifer Lackey for organising this incredible event. Thanks to her and Sam Baron, Kyle Blumberg, Cameron Boulton, Sam Carter, Emmalou Davis, Julien Dutant, Amy Flowerree, Rachel Fraser, Veli Mitova, Gillian Russell, Jason Kwall, Jennifer Lackey, Laurie Paul, and Sarah Wright for discussion at the conference or in preparation for it. An extra special thanks to Sandy Goldberg for his written feedback and comments at the conference.

- Foley, R. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of Belief* (pp. 37–47). Springer.
- Gao, J. (2019). Credal pragmatism. *Philosophical Studies*, 176(6), 1595–1617. <https://doi.org/10.1007/s11098-018-1081-z>
- Gardiner, G. (2018). Evidentialism and Moral Encroachment. In M. Kevin (Ed.), *Believing in Accordance with the Evidence*. Springer Verlag.
- Gibbons, J. (2013). *The Norm of Belief*. Oxford University Press.
- Goodman, J., & Holguín, B. (2023). Thinking and being sure*. *Philosophy and Phenomenological Research*, 106(3), 634–654. <https://doi.org/10.1111/phpr.12876>
- Gordon, R. M. (1990). *The Structure of Emotions: Investigations in Cognitive Philosophy*. Cambridge University Press.
- Greco, D. (2015). How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives*, 29(1), 179–201. <https://doi.org/10.1111/phpe.12059>
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5), 1393–1404. <https://doi.org/10.1007/s11098-015-0553-7>
- Holguín, B. (2022). Thinking, Guessing, and Believing. *Philosophers' Imprint*, 22(0), Article 0. <https://doi.org/10.3998/phimp.2123>
- Horowitz, S. (2019). Accuracy and Educated Guesses. *Oxford Studies in Epistemology*, 6.
- Lee-Stronach, C. (2021). Morality, Uncertainty. *The Philosophical Quarterly*, 71(2), 334–358. <https://doi.org/10.1093/pq/pqaa030>
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, C. (2018). The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology. In K. A.-V. Jeff Dunn (Ed.), *Epistemic Consequentialism* (pp. 23–47). Oxford University Press.
- Littlejohn, C. (2020). Truth, Knowledge, and the Standard of Proof in Criminal Law. *Synthese*, 197(12), 5253–5286. <https://doi.org/10.1007/s11229-017-1608-4>
- Littlejohn, C. (2023). The ranges of reasons and creasons. *Asian Journal of Philosophy*, 2(2), 1–10.
- Littlejohn, C., & Dutant, J. (2020). Justification, Knowledge, and Normality. *Philosophical Studies*, 177(6), 1593–1609.
- Littlejohn, C., & Dutant, J. (2021). Defeaters as Indicators of Ignorance. In M. Simion & J. Brown (Eds.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Littlejohn, C., & Dutant, J. (forthcoming). What is Rational Belief? *Noûs*.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25(6), 205.
- Meylan, A. (2014). Epistemic Emotions: A Natural Kind? *Philosophical Inquiries*, 2(1), 173–190.
- Moss, S. (2021). Pragmatic encroachment and legal proof. *Philosophical Issues*, 31(1), 258–279.
- Otsuka, M. (2015). Risking Life and Limb: How to Discount Harms by Their Improbability. In *Identified versus Statistical Lives: An Interdisciplinary Perspective* (pp. 77–93). Oxford University Press. <https://academic.oup.com/book/9938/chapter-abstract/157273132?redirectedFrom=fulltext>
- Ryan, S. (1991). The preface paradox. *Philosophical Studies*, 64(3), 293–307.
- Ryan, S. (1996). The Epistemic Virtues of Consistency. *Synthese*, 109(2), 121–141. <https://doi.org/10.1007/bf00413765>
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press.
- Sharadin, N. (2018). Epistemic Consequentialism: Haters Gonna Hate. In C. Kyriacou & R. McKenna (Eds.), *Metaepistemology: Realism & Antirealism* (pp. 121–143). Palgrave MacMillan.
- Smith, M. (2021). Against Legal Probabilism. In J. Robson & Z. Hoskins (Eds.), *The Social Epistemology of Legal Trials* (pp. 92–106). Routledge.
- Smith, M. (forthcoming a). Probability, normalcy and the right against risk imposition. *Journal of Ethics and Social Philosophy*.
- Smith, M. (forthcoming b). Risky belief. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12894>

- Sprenger, J., & Hartmann, S. (2019). *Bayesian Philosophy of Science*. Oxford University Press.
- Sturgeon, S. (2008). Reason and the Grain of Belief. *Noûs*, 42(1), 139–165.
- Talbot, B. (2014). Truth promoting non-evidential reasons for belief. *Philosophical Studies*, 168(3), 599–618.
- Taurek, J. (2021). Reply to Parfit's 'Innumerate Ethics'. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (Eds.), *Principles and Persons* (pp. 311–322). Oxford University Press. <https://doi.org/10.1093/oso/9780192893994.003.0015>
- Thomson, J. J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems*, 49(3), 199–219. <https://doi.org/10.2307/1191633>
- Weatherson, B. (2005). Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives*, 19(1), 417–443. <https://doi.org/10.1111/j.1520-8583.2005.00068.x>
- Whiting, D. (2022). *The Range of Reasons: In Ethics and Epistemology*. Oxford University Press.