

- Lange, M. 2022a. Putting explanation back into ‘inference to the best explanation’. *Noûs* 56: 84–109.
- Lange, M. 2022b. In defense of really statistical explanations. *Synthese* 200.
- Leslie, J. 1989. *Universes*. London and New York: Routledge.
- Martin, B. 1998. Coincidences: remarkable or random? *Skeptical Inquirer* 22: 23–8.

Justification, Evidence and Truth

BOB BEDDOR

Being Rational and Being Right

BY JUAN COMESAÑA

Oxford University Press, 2020. 240 pp

1. Introduction

Rational thinkers respect their evidence. This much is a platitude. But when we try to put some flesh on the bones of this platitude, we quickly find ourselves embroiled in difficult questions. What does an agent’s evidence consist in? And how does respecting the evidence relate to justified belief?

Bayesian epistemology offers an elegant framework for modelling rational responses to the evidence. But it leaves these foundational questions unanswered: textbook statements of Bayesianism are usually silent on how to conceive of evidence, or how the rational requirements they espouse link up with justification.

Comesaña’s important book, *Being Rational and Being Right*, seeks to answer these questions. While the book covers a lot of ground, its central contribution is an articulation and defence of a particular theory of evidence:

Experientialism: An agent’s evidence consists in those beliefs which are *ultima facie* justified by experience.

According to Comesaña, experientialism plays nicely with our best decision theory, and it can serve as the backbone for a comprehensive theory of epistemic justification. Rival views of evidence lack these virtues, or so Comesaña argues.

Being Rational and Being Right is filled with fascinating arguments, and it offers an exciting vision of how epistemic and practical rationality fit together. Here, I want to focus on three themes in Comesaña’s extraordinarily rich book: (i) the criticism of factulist accounts of evidence and the arguments for experientialism, (ii) Comesaña’s proposed rule for updating credence in light of new evidence and (iii) the overall theory of justified belief that emerges in the last chapter.

2. A lifeline for factualism?

One of Comesaña's main sparring partners is the factualist, who claims:

Factualism: An agent's evidence consists in various facts.

The most well-known version of factualism is E=K: the view that one's total evidence consists of one's total knowledge (Williamson 2000). But E=K is not the only factualist option. One could also hold that one's evidence consists of a privileged subset of one's knowledge – for example, one's observational knowledge (Maher 1996) or one's epistemic certainties (Beddor 2020). This is important since one of Comesaña's most powerful arguments against E=K concerns the sufficiency of knowledge for evidence (Chapter 4.2). Combining E=K with the idea that rational thinkers conditionalize on their evidence entails that everyone should have credence 1 in everything they know. But in many cases – particularly cases of inductive knowledge – this seems like overconfidence. I am in complete agreement with this point. However, as Comesaña acknowledges, it leaves open the possibility that knowledge – or at least some factive relation – is *necessary* for evidence possession.

2.1 Evidence for factualism?

Before tackling Comesaña's main argument against factualism, let me outline what I take to be the strongest argument on the factualist's behalf. The argument starts by reflecting on the function of evidence. One central function of evidence is to rule out possibilities:

Ruling Out: If S's evidence entails p , then S is entitled to ignore all not- p possibilities for the purposes of practical reasoning.

This much is common ground: Comesaña's decision-theoretic framework in Chapter 4 relies on Ruling Out.

Ruling Out can sound like a truism, but it imposes a substantive constraint on any theory of evidence. In particular, an adequate theory of evidence should explain why Ruling Out holds. Why does possessing p as evidence entitle you to ignore all not- p possibilities? Factualists have a simple answer. If your (factive) evidence entails p , not- p is guaranteed to be false, in which case there is no need to take not- p possibilities into account when deliberating about what to do.

Now, it might seem that non-factualists have an equally good answer. According to many non-factualists (including experientialists), you only possess p as evidence if you *believe* p . Perhaps this explains why Ruling Out holds: if you possess p as evidence, then you are committed to believing that not- p is false, which entitles you to ignore all not- p possibilities.

But this alternative answer faces a worry: it can be rational to take seriously the possibility that some of our beliefs – even some of our experientially justified beliefs – are mistaken. Suppose Deb the detective is investigating a

complex case. After carefully sifting through a byzantine web of evidence – a tangled mess of conflicting alibis, torturous money trails, lots of red herrings, etc. – Deb infers the butler is guilty (call this proposition ‘GUILTY’). As a matter of fact, all of Deb’s beliefs about the case are *ultima facie* justified by her experiences, and these beliefs entail GUILTY. Still, Deb has some doubts about whether she has correctly interpreted and properly weighed all of the relevant considerations. Consequently, she takes further actions – consulting a colleague, gathering still more evidence – before making the arrest. Intuitively, it is rational for Deb to harbour doubts about whether GUILTY is true, and to take further actions allay to these doubts; to simply dismiss the possibility that she got things wrong would manifest an irrational form of overconfidence.¹ But this intuition stands in tension with experientialism. According to experientialism, Deb’s evidence consists of those beliefs which are *ultima facie* justified by her experiences, and we’ve stipulated that these beliefs entail GUILTY. By Ruling Out, it follows that Deb is entitled to dismiss the possibility that the butler is innocent.

Is there any way for experientialists to respond? One option would be to dig in their heels and insist that Deb is entitled to dismiss the possibility that GUILTY is false. It’s just that Deb is not in a position to realize this, because she is not in a position to justifiably believe that all of her beliefs about the case are *ultima facie* justified. However, I suspect that this response would be unwelcome to Comesaña, given the accessibility constraint on the justification that he endorses (to be discussed shortly). Moreover, we can easily amend the case. Suppose a trustworthy oracle tells Deb: ‘Good news! All of your beliefs about the case are *ultima facie* justified’. Even after learning this, wouldn’t it still be rational for Deb to have some doubts about GUILTY? After all, we can suppose that Deb is aware that justification does not entail truth, so the oracle’s announcement still leaves open the very real possibility that she made a mistake. By contrast, if the oracle had told her that each of her beliefs about the case is *true*, she would be entitled to dismiss the possibility that GUILTY is false.

Having offered an argument in favour of factualism over experientialism, let us turn to Comesaña’s main objection.

2.2 *The case against factualism*

Comesaña argues that factualism makes implausible predictions about the following pair of cases:

GOOD LUCAS: Tomás would like to eat some candy. Lucas offers him some, and Tomás reaches for it and puts it in his mouth. Everything goes as planned and Tomás enjoys some candy.

1 This case is essentially a version of the famous preface paradox (Makinson 1965).

BAD LUCAS: Tomás would like to eat some candy. Lucas offers him a marble that looks just like candy, and Tomás reaches for it and puts it in his mouth. Tomás is disappointed.

What is the problem for factualism here? Factualists are committed to saying that the proposition that <That candy-looking thing is candy> (call this proposition ‘CANDY’) is part of Tomás’s evidence in GOOD LUCAS, but not in BAD LUCAS. In some passages, Comesaña claims that this commits factualists to say that Tomás’s action is rational in GOOD LUCAS, but irrational in BAD LUCAS.²

I think this is a bit quick: factualists have some wiggle room here. Even in BAD LUCAS, there are still many other true propositions that could comprise Tomás’s evidence. The factualist can say that these propositions are what make his action rational in BAD LUCAS.

Comesaña anticipates a version of this response and rejects it:

Maybe, if proponents of Factualism are lucky, Tomás [in BAD LUCAS] knows that what Lucas is offering looks like candy. But if he knows that, he also knows that what Lucas is offering looks like a marble that looks like candy. In any case, Tomás doesn’t know enough to rationalize doing what he did. So, according to Factualism, Tomás was irrational. (2)

But perhaps we just need to look further to find the relevant bit of evidence. Here is something else that Tomás knows:

LIKELY CANDY: The candy-looking thing that Lucas is offering is most likely to be candy.

Factualists could propose that LIKELY CANDY is part of Tomás’s evidence in BAD LUCAS, and this is what makes his action rational.³

Still, this point only goes so far. After all, factualists are committed to saying that $\text{Pr}(\text{CANDY}) = 1$ in GOOD LUCAS, but $\text{Pr}(\text{CANDY}) < 1$ in BAD LUCAS. Given this, we should be able to modify the cases so that this probabilistic difference impacts what it is rational for Tomás to do. For example, suppose Tomás has to pay \$1 to get the candy-like thing. And suppose Tomás

2 For example, ‘Knowledge-based decision theory has the consequence that the state where Lucas is offering a marble to Tomás should be part of the decision matrix for Bad Lucas but not for Good Lucas. This means, in turn, that according to knowledge-based decision theory Tomás’s action of tasting what Lucas gave him is rational in Good Lucas but irrational in Bad Lucas’ (Comesaña 2020: 87).

3 Some might worry that this explanation of Tomás’s action is unnecessarily complicated. Does Tomás really need to have knowledge of the probabilities to act rationally? In response, factualists could appeal to the idea that credence can constitute probabilistic knowledge (e.g. Moss 2013, 2018, Beddor and Goldstein 2021). Presumably, all parties will grant that Tomás’s action is guided by a higher credence in CANDY than its negation. Factualists could maintain that if this credence constitutes knowledge, it can be a part of Tomás’s evidence.

values eating an (actual) candy at \$1.01. Then the action that maximizes expected utility conditional on CANDY is *buy the object & put it your mouth*, whereas the action that maximizes expected utility conditional on LIKELY CANDY is *refrain from buying the object*. In this modified pair of cases, the factualist is forced to say that Tomás's action is rational in the good case, but not in the bad case.

So I think Comesaña's basic point is correct. Factualists are committed to saying that the truth-value of a proposition that an agent believes can affect what it is rational for that agent to do. But is this consequence really so unpalatable?

Comesaña thinks so. He argues that this consequence is inconsistent with an accessibility constraint along the following lines:

Accessibility Constraint: If a condition *C* is such that its obtaining makes some attitude or action on the part of a subject irrational, then it must be the case that it is rational for the subject to believe that *C* obtains.
(90)

But should we accept the Accessibility Constraint? Most externalists will reject this principle. Here I don't just have in mind extreme Knowledge Firsters who maintain that justification entails knowledge (like Comesaña, I find such views implausible). Traditional process reliabilists will also reject the Accessibility Constraint. According to process reliabilists, the rationality⁴ of a belief depends on the reliability of the process responsible for that belief. But when an agent's belief is the result of an unreliable process, it is not always rational for them to believe that it is. Of course, some might conclude: 'So much the worse for traditional reliabilism!' Fair enough, my point is simply that whether one finds the Accessibility Constraint plausible will depend on one's broader allegiances in the old externalist/internalist fight. If we are willing to accept externalism about rational belief, embracing externalism about rational action seems like a natural next step.

A second worry is that the Accessibility Constraint is vulnerable to counterexamples. Comesaña anticipates some of these counterexamples, and points out that we may need to restrict the principle to cases where the condition is 'non-epistemic'. As Comesaña notes, this raises the tricky task of saying what makes a condition qualify as 'epistemic'. But even if this can be done, a restricted accessibility constraint will still need to deal with a further batch of counterexamples. Consider cases where an agent does not have conceptual resources to understand condition *C*. For example, the rationality of a small child's perceptual belief may depend in part of the condition that *lighting conditions are normal*. But the child may lack the concept of normal

4 Really justification, but for present purposes it is harmless to assume we can use these normative statuses interchangeably.

lighting conditions, and so may not be in a position to rationally believe this condition obtains. Or consider cases where parts of my mental life are hidden from me. For example, suppose I have a tacit belief q , which is not introspectively available to me (it would only be revealed through psychotherapy, or by being offered certain bets that I have never contemplated). Suppose my belief in q is inconsistent with one of my explicit beliefs, p . Arguably, the fact that I believe q makes it irrational for me to believe p , or at least diminishes the rationality of the latter belief. But it does not follow that it is rational for me to believe that I believe q – after all, I have no introspective access to my belief in q .

Comesaña tries to ward off this style of attack, writing: ‘Merely hurling counterexamples at the principle does not amount to a defense of knowledge-based decision theory from my argument’ (2020: 91). But I wonder whether the counterexample counterattack can be countered so quickly. Sometimes a steady stream of counterexamples is an indication that a principle is off-track altogether, rather than in need of further Chisholming.

Taking stock: Comesaña’s argument against factualism is certainly important; it goes right to the heart of factualism’s commitments. However, I’ve argued that these commitments are not as problematic as they initially appear to be. They will certainly look problematic to those who are committed to the Accessibility Constraint. But there are independent reasons to question whether any version of the Accessibility Constraint is plausible. Now, all these defensive manoeuvres might seem ad hoc if there was not a powerful argument in favour of factualism. But I’ve suggested that there is: factualism is able to explain why Ruling Out holds, whereas experientialism has trouble doing so.

3. *Ur-Prior Conditionalization*

Another important contribution of *Being Rational and Being Right* is Comesaña’s defence of a particular updating rule: ur-prior Conditionalization. According to Ur-Prior Conditionalization, your credence in p at some time t should equal the conditional probability that your ur-prior (represented ‘ Cr_u ’) assigns to p given your total evidence at t :

$$\text{Ur-Prior Conditionalization: } Cr_t(p) = Cr_u(p|e_t)$$

One of Comesaña’s main arguments for Ur-Prior Conditionalization is that it avoids the ‘stickiness’ of extremal probabilities under other updating rules, such as (standard) conditionalization. To illustrate, Comesaña offers a helpful example:

Let us suppose that I know that the color of the wall (red or white) is going to be decided by the flip of a fair coin, and the color of the lights (also red or white) is also going to be decided by a second flip of

the same fair coin. If we let K represent all this background information, then, conditional on K , the ur-prior assigns the same probability (namely, $\frac{1}{4}$) to each of the state descriptions determined by the partition $\{E, \neg E, D, \neg D\}$, where E is the proposition that the wall is red and D is the proposition that the lights are red. ... When I open my eyes and look at the wall, I have an experience with the content that E and, as a result, I have E as evidence. Therefore, letting Cr_1 be my credence function at this time, $Cr_1(E|D)=Cr_1(E)$. This might seem as the wrong result, for we want my credence that the wall is red to be affected by learning that the lights are red. But my credence will be affected when I later learn that the lights are red, for, letting Cr_2 be my credence function at that time, then $Cr_2(E)=Cr_2(E|D)<Cr_1(E|D)$. (Comesaña 2020: 145)

While this is an ingenious treatment of defeat, I want to raise a concern. Consider my situation at t_1 , when I am looking at the red wall, before I have learned that the lighting is red (D). At t_1 , I am certain the wall is red (E), and this is part of my evidence. So I should think that if I were to learn D at t_2 and lower my credence in E , my resulting credence in E would be less accurate than my current credence. At the same time, if I know for sure that Ur-Prior Conditionalization is the right updating rule, I know that if I learn D at t_2 I will be rationally required to lower my credence in E . I thus seem to be in a dilemma. On the one hand, I anticipate that lowering my credence in E in response to D will be an epistemic loss (loss of accuracy); on the other hand, I realize that not lowering my credence in E in response to D will be an epistemic sin (a violation of the updating rule). What should I do in this predicament? This is unclear. One option is to go the dogmatist's route: stick my head in the sand, and guard against the possibility of ever learning D . But, intuitively, such behaviour would be highly irrational. Note that this issue does not arise if we update by conditionalization, due to the very feature that Comesaña deems a bug: since credence 1 is sticky, learning D won't reduce my credence in E , so I have nothing to fear.

Having raised this problem, let me sketch a possible solution. One option would be to combine Ur-Prior Conditionalization with a 'shifty' view of evidence, according to which an agent's evidence shifts with features of their context.⁵ Now, Comesaña is already committed to some form of evidential shiftiness: on his view, my evidence at t_1 includes E , but at t_2 I lose E as evidence. Perhaps we should go a step further, and hold that an agent's evidence shifts depending on which possibilities they are taking seriously (or perhaps which possibilities they *ought* to be taken seriously).⁶

5 This could be implemented in either a subject sensitive invariantist or a contextualist setting, though in the latter case it should be formulated in terms of the shiftiness of 'evidence' talk, rather than evidence itself. Cf. Greco (2017), which defends a contextualist theory of evidence.

6 This is analogous to Lewis's (1996) 'Rule of Attention' governing knowledge ascriptions.

How does this help? As soon as I begin to consider the possibility that I might learn D , this very act of contemplating this possibility affects my evidence. I now no longer have E as evidence, and hence my credence in E is no longer 1. Since my credence in E is no longer 1, the prospect of learning D no longer fills me with alethic dread, since from my new perspective, it will not guarantee a loss in accuracy. On this view, receiving new evidence is not the only thing that can rob me of evidence; the very prospect of receiving new evidence can trigger this evidential loss. In principle, this shifty manoeuvre could be incorporated into any number of theories of evidence, including both factualism and experientialism. Factualists might try to derive the shiftiness of evidence from the shiftiness of knowledge or certainty; experientialists could try to derive the shiftiness of evidence from the shiftiness of justification.⁷

4. Comesaña's account of justification

Let me turn now to Comesaña's positive account of justification. Comesaña develops a nuanced view that weaves together evidentialist and reliabilist elements, offering a successor to his earlier evidentialist reliabilist hybrid (2010). The updated account goes like this:

Coarse-Grained Evidentialist Reliabilism: A belief that p by S is justified if and only if:

Either:

1. S 's experiences provide him with p ; or
- 2a. S 's experiences provide him with e ;
- 2b. the belief that p by S is based on e ;
- 2c. $Cr_u(p|e) \geq r$;
- 2d. There is no more inclusive body of evidence e' had by S such that $Cr_u(p|e') < r$.⁸

⁷ An interesting question is whether a shifty view of justification could help address my earlier objection to experientialism involving Ruling Out. Perhaps, some might suggest, Deb's evidence initially entails GUILTY. But once Deb starts considering the possibility that she made a mistake, her evidence shifts: beliefs that used to be *ultima facie* justified lose this status, and so her evidence no longer entails GUILTY (cf. Clarke (2017), who defends a 'sensitivist' response to the preface paradox along these lines). This is an intriguing view, and may offer the best response to the worry. That said, it raises some important questions and concerns. For example, which of her previously justified beliefs about the case lose their justificatory status? It might be that there is no particular belief that she deems particularly suspect. Perhaps they *all* cease to qualify as justified. But one might well worry that this is an overly sceptical conclusion. A related concern is that this response entails that the oracle's pronouncement ('Good news! All of your beliefs about the case are *ultima facie* justified') must be false; indeed, the mere fact that Deb is doubting the conjunction of her beliefs is enough to refute the oracle's proclamation. But, on the face of it, it seems like what the oracle says could well be true. An interesting question for further research is whether these questions and concerns can be satisfactorily dealt with.

⁸ As Comesaña notes, there is a straightforward way of extending this account to provide a theory of justified credence.

Comesaña's view is formulated with admirable rigour and clarity, and I am very sympathetic to the spirit of his reliabilist/evidentialist hybrid. That said, I want to raise two potential concerns.⁹

4.1 Renouncing naturalistic aspirations

Historically, reliabilists have wanted to explain justification in purely naturalistic terms, without relying on any unreduced epistemic notions. This goal is front and centre in Goldman's pioneering work; just consider his claim, 'I want a theory of justified belief to specify in non-epistemic terms when a belief is justified' (1979: 90). In my eyes, this is the main source of reliabilism's explanatory appeal. The reliabilist offers to write a chapter in the naturalist's book of the world – a book that tells us how all normative properties link up to non-normative properties.

Comesaña rejects this naturalistic aspiration. According to him, Cr_u is an evidential probability function, not explicable in terms of objective probabilities. Now, I agree that if we were to construe Cr_u in terms of objective probability, Evidentialist Reliabilism would be subject to counterexample. Still, there is a worry that the view avoids counterexample by abandoning one of the features that made traditional reliabilism appealing.

Some may think this is too quick; Evidentialist Reliabilism is compatible with some versions of naturalism. In particular, suppose we characterize naturalism as a supervenience thesis, according to which justificatory facts supervene on natural facts. Evidentialist Reliabilism is perfectly consistent with this supervenience claim.¹⁰

However, we should distinguish between (i) a view that is *consistent* with the naturalist's supervenience thesis, (ii) a view that *explains* why the naturalist's supervenience thesis holds. Traditional reliabilism is a thesis of the second type. As Kim (1988) notes, one advantage of a view that explicitly identifies justification with some particular natural property (as traditional reliabilism does), is that it explains why justificatory facts supervene on natural facts. After all, if the property of being justified just is the property of being reliably formed, then there is no great mystery as to why fixing the natural facts will fix all of the justificatory facts. By contrast, Evidentialist Reliabilism is a view

9 First, one small point: clause (2d) runs into trouble when S has a more inclusive body of evidence e' such that $Cr_u(p|e') < r$, but S also has a still more inclusive body of evidence e'' such that the $Cr_u(p|e'') > r$. Evidentialist Reliabilism predicts that S's belief is not *ultima facie* justified, which is intuitively the wrong result. A simple fix would be to let e_{total} denote S's total evidence, and replace (2d) with the requirement that $Cr_u(p|e_{\text{total}}) \geq r$. See Pettigrew 2021.

10 Thanks to Juan Comesaña for discussion here.

of the first type: it is consistent with the supervenience of the epistemic on the natural, but it does not explain why this supervenience holds.

4.2 *Is the account predictive?*

A related concern is whether Evidentialist Reliabilism yields predictions about whether a subject is justified in believing a proposition in a given situation. If we understood the *ur*-prior in terms of objective probability, we would be able to use our independent purchase on objective probability (insofar as we have one) to generate predictions about specific cases. But, as we've seen, Comesaña understands the *ur*-prior in terms of evidential probability. Now, perhaps we could use our intuitions about evidential probability to guide us. But are intuitions about evidential probability sufficiently distinct from intuitions about justification? Some might worry that when we make an intuitive assessment about whether some body of evidence *e* provides evidential support for *p*, what we are really doing is considering whether *e* justifies one in believing *p*. If so, we won't be able to test the theory against our pre-theoretic intuitions about justification, since we will be relying on these pre-theoretic intuitions to tell us what the theory predicts.

Here too, it is helpful to contrast Evidentialist Reliabilism with more traditional reliabilist views. Traditional process reliabilist views helped themselves to a notion of reliability that was defined independently of any epistemic terms (say, in terms of the ratio of true to false beliefs that the process produces). Process reliabilists then sought to support their view by establishing a close correlation between this independently defined notion of reliability and our intuitive verdicts about justification. For example, beliefs based on wishful thinking, guesswork or hasty generalization are intuitively unjustified; these belief-forming processes are also unreliable. Similarly, beliefs based on standard perceptual processes in normal conditions, or competent deduction from logical axioms are intuitively justified; these belief-forming processes are also reliable (Goldman 1979: 95). Of course, subsequent philosophers have questioned how close of a correlation there really is; here the well-known counterexamples to traditional reliabilism rear their many heads. My point is not that traditional process reliabilism is correct. Rather, it is that traditional process reliabilism generated clear predictions about whether someone is justified – indeed, it is precisely this feature of traditional process reliabilism that rendered it susceptible to counterexample.

4.3 *Taking stock*

Comesaña's Evidentialist Reliabilism offers to join together two major epistemological traditions: reliabilism and evidentialism. While this is an exciting union, I've suggested that it comes with some costs. How heavy are these costs? Should they lead us to retreat to a 'pure' reliabilist theory, or to

embrace an alternative hybrid view?¹¹ One's answers to these questions will depend, in part, on one's general philosophical proclivities – for example, whether one longs for a naturalistic reduction, and whether one wants one's epistemological theories to predict our intuitions rather than accommodate them.

National University of Singapore
Singapore
rbeddor@gmail.com

References

- Alston, W. 1988. An internalist externalism. *Synthese* 74: 265–83.
- Beddor, B. 2020. New work for certainty. *Philosophers' Imprint* 20: 1–25.
- Beddor, B. 2021. Reasons for reliabilism. In *Reasons, Justification, and Defeat*, eds. J. Brown and M. Simion, 146–76. Oxford: Oxford University Press.
- Beddor, B. and S. Goldstein. 2021. Mighty knowledge. *Journal of Philosophy* 118: 229–69.
- Clarke, R. 2017. Preface writers are consistent. *Pacific Philosophical Quarterly* 98: 362–81.
- Comesaña, J. 2010. Evidentialist reliabilism. *Noûs* 44: 571–600.
- Comesaña, J. 2020. *Being Rational and Being Right*. Oxford: Oxford University Press.
- Goldman, A. 1979. What is justified belief? In *Justification and Knowledge*, ed. G. Pappas, 143–64. Dordrecht: Reidel.
- Goldman, A. 2011. Towards a synthesis of reliabilism and evidentialism. In *Evidentialism and Its Discontents*, ed. T. Dougherty, 254–80. Oxford: Oxford University Press.
- Greco, D. 2017. Cognitive mobile homes. *Mind* 126: 93–121.
- Kim, J. 1988. What is 'naturalized epistemology'? *Philosophical Perspectives* 2: 381–405.
- Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74: 549–67.
- Maher, P. 1996. Subjective and objective confirmation. *Philosophy of Science* 63: 149–74.
- Makinson, D.C. 1965. The paradox of the preface. *Analysis* 25: 205–7.
- Miller, E. 2019. Liars, tigers, and bearers of bad news, oh my! Towards a reasons account of defeat. *The Philosophical Quarterly* 69: 82–99.
- Moss, S. 2013. Epistemology formalized. *Philosophical Review* 122: 1–43.
- Moss, S. 2018. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Pettigrew, R. 2021. What is justified credence? *Episteme* 18: 16–30.
- Tang, W.H. 2016. Reliability theories of justified credence. *Mind* 125: 63–94.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

11 For other views that combine reliabilist and evidentialist elements, see, for example, [Alston 1988](#), [Goldman 2011](#), [Tang 2016](#), [Miller 2019](#), [Beddor 2021](#) and [Pettigrew 2021](#).