

Knowledge and Prizes

Clayton Littlejohn (Dianoia Institute of Philosophy)

Julien Dutant (King's College London)

Forthcoming in A. Logis and Jacques-Henri Vollet (ed.) *Putting Knowledge to Work: New Directions for Knowledge-First Epistemology*. Oxford University Press.

Abstract

We examine two leading theories of rational belief, the *Lockean* view and the *explanationist* view. The first is appealing because it fits with some independently plausible claims about the ways that rational persons pursue their aims. The second is appealing because it seems to account for intuitions that cause trouble for the Lockean view. While fitting the intuitive data is desirable, we are troubled that the explanationist view seems to clash with our theoretical beliefs about what rationality must be like. We think that upon further examination, the intuitive appeal of the explanationist view starts to diminish. We also think that these further intuitions that spell trouble for the explanationist spell trouble for any theory that is not *expectationist*. We propose a novel expectationist theory of rational belief that improves upon the Lockean and the explanationist views. We think that recent defences of the Lockean view contain an important insight. A substantive theory of rational response should be based on a suitable theory of prizes and a suitable theory of how we should pursue prizes in the face of uncertainty. Most theories of rational belief typically take for granted a truth-centred picture of epistemic prizes (e.g., that epistemic desirability and undesirability can be fully understood in terms of accuracy) and then differ in terms of how they recommend pursuing prizes so understood. We think the Lockean embrace plausible principles of how prizes should be pursued. We trace the difficulties that this view faces to veritistic assumptions about prizes. We suggest that some prizes are *epistemically loaded* in that a complete description of the prize will itself make reference to our epistemic states or standards. We argue that knowledge matters to rational belief and choice because in the epistemic domain, knowledge is the prize. We see this in practical domains, too. In some choice settings, what's desired is desired, in part, because it involves a kind of connection to reality only knowledge provides.

0. Introduction

Many epistemologists endorse a broadly evidentialist approach to epistemic rationality.¹ They might say that a belief is rational iff it 'fits' the evidence or the

¹ We use 'evidentialism' in the way that Conee & Feldman (2004) and McCain (2014) do. It is the view that a thinker's evidence at a time entirely determines what is (*ex ante*) rational for her to believe at that time. Sometimes the term 'evidentialism' is associated with the view that every epistemic reason should be thought of in terms of evidence. (See, for example, Shah (2006).) Some authors (e.g., Owens (2000)) have argued that some epistemic reasons (e.g., the reasons a thinker has *not* to believe) are

evidence provides 'sufficient' support for that belief. These claims can seem truistic. Things get interesting when we try to say what fit or sufficiency amounts to.

In this paper, we will evaluate two familiar theories of fit and offer a third. According to the first, beliefs 'fit' the evidence when the evidence provides sufficiently strong support for them (and fails to fit the evidence otherwise). Given plausible assumptions about the relationship between strength of evidential support and rational degrees of belief, this is a veritistic *strength-centred* view. It is similar to the Lockean view of rational belief (Dorst (2019), Easwaran (2016), Foley (2009), Sturgeon (2008)). On this view, it's always rational to believe if the probability of the target proposition on the thinker's evidence is sufficiently high (High) and it's not rational to believe if the probability of that proposition is insufficiently high (Low).² According to the second, we should think of rational support in terms of explanation (McCain (2014, 2015), McCain and Moretti (2021)).³ On this *explanationist* view, it is rational for a thinker to believe at a time when that belief stands in the right explanatory relation to the evidence the thinker has at that time.

We think there are *prima facie* plausible arguments for the veritist strength-centred view as well as the explanationist view. If we're right that they disagree about a key case, this is a diplomatic way of stating that there are *prima facie* plausible *objections* to each of these views. We see our proposal as offering a way to build on these approaches. We incorporate aspects of the strength-centred theory and find a role for explanation in the theory of rational belief in a way that overcomes the difficulties that these two more familiar approaches face.

Just to put our cards on the table, we think that the difficulties that arise for these views, has two sources. We think that our theories of rational belief should tell us something about epistemic *prizes* and something about the proper ways to *pursue* prizes. A theory of epistemic prizes tells us what features of our beliefs might make them objectively desirable or undesirable from the epistemic point of view. It further should tell us something about how desirable or undesirable happy or unhappy results might be. A theory of rational pursuit tells us which responses are rational

not necessarily things we should identify with evidence even if it's true that all epistemic reasons supervene upon a thinker's evidence. Nothing we say here seems to turn on whether evidentialism in this second sense is true.

² Appley and Stoutenburg (2017) seem to assume something like High in their critical discussion of explanationism. We see rejecting High as a wise move for the explanationist to try to undercut their objections, but there might be other strategies available to the explanationist. (As we'll see below, rejecting High might come with costs, too.)

³ We focus on the work of McCain and Moretti (2021) for the most part for two reasons. First, they provide the most recent defences of the explanationist approach. Second, we note that it's often unclear whether the explanationist is interested in credence, confirmation, or outright belief. We're primarily interested in outright belief and the norms that govern this notion of belief. The same holds true for McCain and Moretti. We do not discuss explanationist work on credal norms or confirmation.

given our information and assumptions about prizes. When we find putative counterexamples to the Lockean view, the assumption that beliefs are objectively epistemically desirable iff accurate is usually not questioned. The first move is to find some relation between the evidence and truth that differs from the assumption proposal that it's rational to pursue these desirable states by forming those attitudes that are sufficiently likely to have the properties we desire. This is a move we shouldn't make. We should critically examine the veritistic assumptions about prizes.

1. Two perspectives on strength

In this section, we'll present two arguments. The first purports to show that we should reject High. The second purports to show that we should accept High. We think it's important for anyone offering a theory of rational belief to say something in response to these arguments.

1.1 The Lottery Argument

Much has been made in the literature about claims like these:

Hearing: if someone tells you that p , you might, knowing little else, come to rationally believe that p .

Seeing: if it seems visually as if p , you might, knowing little else, come to rationally believe p .

Playing: if you hold a ticket for a fair lottery, you will not, knowing little else, come to rationally believe that it lost.⁴

Not everyone accepts these claims, of course, but we are satisfied with the things people have said in support of them. We will look for ways to explaining Hearing, Seeing, and Playing rather than ways of trying to explain away the intuitions that underwrite them.

Getting the precise details of what's going on in hearing and seeing cases that confers rational support on your beliefs is tricky, but the idea is that testimony and sense experience can provide justification for beliefs when the kinds of grounds

⁴ Note that we said, 'little', but not 'nothing'. On the explanationist view we'll consider, testimonial beliefs and experiential beliefs are rational because of an explanatory inference and that might require more than just knowledge of how things sound or look, say. Intuitions about lottery-type cases have loomed large in the arguments for knowledge-centred theories of rational belief. See, for example, Bird (2007), Hirvelä (2022), Ichikawa (2014), Kelp (2014), Littlejohn & Dutant (2021), Rosenkranz (2021), Sutton (2005), and Williamson (2000). We should note that intuitions about the lottery seem to be relatively robust, as evidenced by the robust intuitions that people have about statistical evidence cases in the law. See Gardiner (2019) and Moss (2018) for discussion. We agree with Smith (2016) and others that it is awkward at best to say that jurors can rationally be convinced of a defendant's guilt on the basis of naked statistical evidence but shouldn't convict. Perhaps this is because, as Conee (2004) has suggested, a natural way of understanding justified or rational belief is in terms of a reasonable doubt standard.

we have in a lottery case (when the only information we happen to have about the outcome is extracted from our knowledge of the set up and we know the lottery to be fair) does not make it rational to believe outright (Harman (1968)).

It's clear that this sort of contrast causes trouble for the veritist's strength-centred view. Consider three further claims:

Hearing+: if someone tells you that p , you might, knowing little else, come to rationally believe that p and rationally be very confident in p .

Seeing+: if it seems visually as if p , you might, knowing little else, come to rationally believe p and be rationally very confident in p .

Playing+: if you hold a ticket for a fair lottery, you will not, knowing little else, come to rationally believe that it lost but you will rationally be nearly certain that you lost.

Rationality attaches to the experientially grounded and testimonially supported beliefs but not to lottery beliefs even when you should be more confident in the lottery belief than, say, the belief that Arsenal lost (which you acquire by reading the paper) or the belief that it's raining outside (which you acquire by seeing the falling water outside your window). The probability of a mistake in the paper is small, but it's not as small as the probability of winning the lottery.

It's difficult to make sense of these intuitions if we just think of rational support in terms of strength.⁵ The intuitions that underwrite these claims are an obvious threat to High. If we think of sufficient strength as allowing for justification in the testimony or sense experience case, this seems to let the lottery beliefs in. If we want to keep the lottery beliefs out and raise the bar accordingly, we have to concede that many of our testimonial and experiential beliefs are not rational.

If we reject the strength-centred theory that connects rational degrees of belief to rational belief on these grounds, we'll have to allow for this initially surprising possibility – that there might be pairs of propositions, p and q , where we should be more confident of the former but can only rationally believe the latter. This is the price we must willingly pay to make sense of our intuitions that seem to vindicate Hearing, Seeing, Playing, and their strengthened counterparts.⁶

Nelkin (2000) and Smith (2016) note that there is this important difference between the testimonial and experiential cases and the lottery case. While we should expect more errors in the testimonial case and experiential case than in the lottery case, we don't respond the same way to the discovery of such errors. When we discover that the target proposition in the lottery case is false (and calm down a bit having discovered that we've just become very rich), we don't think that this calls for any sort of special explanation. When, however, we discover that the target

⁵ We say 'difficult', but not impossible. See below.

⁶ This point has convinced some authors that belief could not be understood in terms of having a high degree of confidence. See Jackson (2019) and Littlejohn (2015) for discussion of this case. See Leitgeb (2014) for an interesting positive proposal about how belief and credence might be linked that delivers the same (negative) verdict about lottery cases.

proposition is wrong in the testimonial case or the experiential case, we do think that this calls for a kind of explanation. This seems to be an important contrast between the cases. When the grounds or evidence makes it rational to believe, it might not guarantee that the target proposition is true or show that it's more likely to be true than some lottery proposition, but there seems to be some sort of explanatory tie between evidence and belief as evidenced by the desire to explain the lack of connection when we discover that the target proposition is false.

Enter the explanationist theory of rational belief (McCain and Moretti 2021).⁷ The explanationist proposes that it is the explanatory connection between the evidence and a target proposition that might make the latter rational for a thinker to believe. This is what fit consists in. If it happens to be the case that some propositions aren't rational to believe despite the fact that they're more likely to be true than some other propositions that are rational to believe, the explanationist will say that this is because high probability is not a sufficient condition for explanatory connection. We already know that and the puzzling pattern of intuitions seems to be neatly accounted for.

We'll focus on this version of explanationism:

Believing p is justified for S at t if and only if at t : (1) S has total evidence, E ; (2) either (i) p is the best (sufficiently good) explanation of e (where e is a subset of E), or (ii) p is an explanatory consequence of the best (sufficiently good) explanation of e (i.e., the relevant explanation of e would provide an explanation of p 's truth that is significantly better than the explanation it would provide of $\sim p$'s truth); (3) it is not the case that p fails to satisfy (i) and (ii) with respect to e because of the additional evidence included in E (McCain and Moretti 2021: 86).⁸

Let's note a few things. First, while we understand why some authors would want to deny that rationality and justification really amount to the same thing, we do not

⁷ There is a long tradition of broadly explanationist views in epistemology. A historically important discussion of inference to best explanation can be found in Russell's (1912: 22) discussion of scepticism. Many explanationists defend the view that inference to best explanation is fundamental in the sense that it's rational force should not be understood in other terms. See Lipton (2001), Lehrer (1970), Lutz (2020), and Poston (2011) for sympathetic discussion. More ambitious is the idea that explanatory inference is somehow fundamental to rational inference as such. In addition to McCain and Moretti (2021), see Lycan (1988). We focus on the most ambitious explanationist views because it seems only this most ambitious version of the view predict that high probability wouldn't be sufficient for rational support in the absence of explanatory connection.

⁸ As Dellsén (2021) observes, many explanationists see their proposal as a proposal about the processes by which rational thinkers assign (subjective) probabilities to hypotheses (e.g., Henderson (2014)). Such views don't directly bear on the questions that primarily interest us, questions about rational binary, full, or outright belief.

distinguish them here. Those who do distinguish them should take our remarks as applying in the first instance to rationality. Second, McCain and Moretti take E to be constituted by some set of mental states. None of our objections target this conception of evidence. Third, we assume that better explanations are better because they are lovelier explanations and exhibit explanatory virtues better than rivals.⁹ Fourth and finally, we take the explanationist view to be primarily a view about what makes it (*ex ante*) rational to believe. We set aside questions about the role that loveliness plays in the rational assignment of subjective probabilities.¹⁰

Here's the key to explaining Seeing, Hearing, and Playing in the explanationist framework.¹¹ Given the setup of the lottery, neither the hypothesis that you'll lose nor the hypothesis that you'll win provides a better explanation of the evidence but they would compete as explanations if offered. If neither does better than the other, neither explanation would be good enough. And if neither explanatory inference is good enough, neither the belief that the ticket lost nor the belief that it won would fit the evidence. Suspension would be the only remaining option.

1.2 The Expectationist Argument

The argument against High is case driven. The argument for High is more theory driven. Consider two *veritist* theories. The first is a theory of epistemically desirable outcomes or prizes. It's inspired by some of William James's (regrettably woolly) remarks about what matters to the would-be believer:

Evaluative veritism: The most desirable outcome (epistemically speaking) is that our beliefs are accurate. The least desirable outcome (epistemically speaking) is that our beliefs are inaccurate. It is more desirable (epistemically speaking) to suspend than to believe falsehoods.

This gives us an ordering in terms of a kind of epistemic value that we can think of as *objective* epistemic value. As James rightly observed, the implications of this value theory for questions about what we should believe (if these are understood in some more *subjective* or *perspectival* way) are not straightforward. If we should attach some great disvalue to believing falsehoods, this will make it harder to find evidence that would make it rational to believe. If, however, we do not attach some great disvalue to believing falsehoods, rational belief might be easier to come by.

⁹ For a fuller discussion of explanatory virtues, see

¹⁰ For concerns about explanationist views that concern subjective probabilities, see Climenhaga (2017).

¹¹ We do not know what our explanationists think about the lottery case. It is not discussed in McCain (2014) or McCain & Moretti (2021), though the idea that the evidence that 'merely' makes it very probable that something is true is not sufficient to warrant outright belief has figured prominently in debates for thinking that something akin to an explanatory connection between evidence and belief is necessary for rationality (e.g., in (Harman (1968), Nelkin (2000), and Smith (2016)).

This way of thinking about the matter seems to rely implicitly on a kind of *expectationist* outlook. It seems to assume that a rational thinker's doxastic 'choices' between the options of belief, disbelief, and suspension will be sensitive to two things: the objective values that would be realised by the possible outcomes and the twin risks of believing falsehoods and failing to believe truths. The expectationist thinks that *if* we rationally should be guided by the evaluative considerations that figure in the veritist view, we should believe (suspend/disbelieve) when believing (suspending/disbelieving) does better in terms of expected veritistic value than the alternatives. While we can debate the merits of views on which the value that's realised by our attitudes depends upon practical factors, whether some epistemologists' instincts reflect too much or too little aversion to believing falsehoods, etc., these can be seen as 'in house' disagreements between people who accept High and Low. Given plausible weightings, it's not at all clear how we could end up rejecting High and embracing claims like Seeing, Hearing, and Playing.

Someone who accepts this sort of outlook might say that the explanationist view is mistaken if it tells us that we should refrain from believing in lottery cases. They might say something like this. Explanatory inferences are fine and good as *instruments* for acquiring true beliefs or rooting out false ones, but the presence or absence of explanatory connections between evidence and belief only matters to the estimation of *prizes* if we happen to have the interests of someone who is curious about and desires beliefs about such explanatory connections. If, however, we remember that the perspective of the would-be believer is that of someone who desires to get hold of the truth and is averse to the acquisition of false beliefs, we have to remember that from *this* perspective the 'decision' to withhold or suspend when the probability of a target proposition is sufficiently high is not reasonable. To suspend in this instance would be to prefer one response (suspension) to another (belief) when it should be evident to the thinker that the favoured response does *worse* in terms of expected desirability.

If the explanationist concedes that high probability without the right explanatory connection is sufficient for rational belief, they might avoid this line of objection, but then they lose the support of the lottery argument.¹² If they wish to reject High, their options are more limited. They could deny that considerations of epistemic desirability and expected epistemic desirability matter to the 'choice' between options. Alternatively, they could say that evaluative considerations do matter but challenge the veritist theory of epistemic value on the grounds that it doesn't treat explanatory loveliness as part of what determines the value of epistemic prizes.

¹² This would not be the only problem that this would create for their view. They would also need to explain how probability connects to loveliness. See Lipton (2001). Would the idea be that an explanation is lovelier than a rival by virtue of its probability on the evidence or the other way around? The explanationist should not use comparative probability to determine comparative loveliness, but then it's not clear how they could argue that the hypothesis that a ticket loses provides a better explanation than the hypothesis that it won.

We see two potential problems with rejecting evaluative veritism on the grounds that it doesn't recognise explanatory loveliness as something that helps determine the value of epistemic prizes. Dialectically speaking, what the expectationist would need to say is something stronger than the claim that being part of the best explanation confers additional value upon a belief. Remember that responding to the lottery argument requires them to say that suspending does better in terms of expected epistemic value than believing in the lottery case in spite of the fact that the target proposition is nearly certain to be true. To explain why suspension would do better than belief in terms of expected epistemic value, they would have to say that it is *undesirable* to believe truths that lacked the right explanatory properties. This seems to us to be a difficult claim to sell.

Moreover, this position on epistemic desirability might actually be in tension with the explanationist view under consideration. Consider clause (ii) in McCain and Moretti's account. They use this to try to explain how things like the logical consequences of propositions justified by clause (i) would be justified. On its face, it seems that some such logical consequences would be *mere* consequences of the best explanations and not themselves things that get positive status by virtue of doing explanatory work. They are, nevertheless, things that the explanationist rightly regards as rationally believed. It thus seems unlikely that they would embrace the expectationist approach in linking rational belief to beliefs that maximise expected epistemic desirability whilst adopting this alternative value theory.

We do not know how the explanationist would respond to the expectationist argument. Nevertheless, we think that it has significant force given certain background assumptions that strike us as plausible (i.e., that rational believers care about expected epistemic desirability and that the veritist theory is at least a plausible approach to thinking about the value of epistemic prizes). We'll see below that there are cases where something akin to the expectationist approach will be needed to make sense of some intuitions about cases. We don't see how to make the explanationist view say the things that the strength-centred view says whilst making sense of the intuitions that support Seeing, Hearing, and Playing, so we hope readers will agree that our view enjoys at least one advantage over the explanationist view.

2. Metacoherence and rational support

In this section, we'll argue that the explanationist view doesn't explain intuitions that a strength-centred theory explains straightforwardly. Before we present the problematic cases, let's discuss a challenge that arises for any theory of rational belief. This is the challenge of accounting for metacoherence constraints. These constraints can be thought of as rational constraints that hold between levels (i.e., constraints that reveal something about the support provided for attitudes concerning first-order propositions and higher-order propositions about these first-order attitudes and the support they receive). Here is a relatively uncontroversial constraint. It is a *local* constraint that concerns *accuracy*: it is not rationally co-tenable to believe *p* whilst believing that this belief is inaccurate or incorrect. If it were rational to hold the beliefs that violated this constraint, it could be rational, say, to

both believe that it's raining and believe that your belief about the rain is mistaken. This, in turn, is to hold beliefs that are inconsistent and it's hard to see how both propositions might enjoy adequate support on the explanationist view. It is a credit to this view that it explains this constraint.¹³

The metacoherence constraints that we usually focus on are more interesting than this one. McCain and Moretti (2021: 133) argue that their explanationist view can vindicate additional local metacoherence constraints. They maintain that their view shows that the evidence cannot simultaneously support (a) believing p and (b) believing that the evidence doesn't support p . In turn, this suggests that they can explain why the evidence cannot simultaneously support (a) believing p and (b) believing that it's not rational to believe p . In both cases, they claim that insofar as p being true is part of the best explanation of the thinker's evidence, neither the lack of evidential support nor the irrationality of the attitude would also provide the best explanation of the thinker's evidence. If we wanted to generalise this to cover the case of knowledge, we might be able to argue that the evidence wouldn't support both (a) believing p and (b) believing p isn't known on the explanationist view. We struggle to think of a case in which p and the subject's failing to know p are both parts of the best explanation of something. We don't know what line the explanationist wants to take on risk, but we might also consider whether the evidence can simultaneously support (a) believing p and (b) believing it's not likely that p . Intuitively, it seems hard to believe that it might be rational to both believe p and it's unlikely that p . On its face, it's hard to imagine a case in which the (apparent) fact that it's unlikely that p and the (apparent) fact that p both figure in the best explanation of something.

The metacoherence constraints that McCain and Moretti discuss are *local* and concern *binary* attitudes. They are local because they concern particular attitudes. They concern binary attitudes like outright belief rather than, say, belief-credence pairs. We think it's important to consider less localised constraints and constraints that hold between beliefs and credences. Let's consider the contrast between two global constraints:

Local accuracy constraint: It is not rationally co-tenable to believe p and believe your belief about p is inaccurate.

Global accuracy constraint: It is not rationally co-tenable to believe what you do and believe that there is one belief you hold that is inaccurate.¹⁴

¹³ To see the inconsistency, we'll assume that the subject has self-knowledge so that she knows that her belief about the rain is the belief that it is raining. If she knows that she believes that it's raining and she believes this content is mistaken, she believes both that it's raining and that it's not true that it's raining.

¹⁴ For defences of this constraint (under different names), see Evnine (1999), Leitgeb (2014), Pollock (1986), Ryan (1991), and Smith (2016). For arguments against accepting this constraint, see Easwaran (2016), Foley (2009), Littlejohn & Dutant (2020), Makinson (1965), Praolini (2019), and Worsnip (2016).

While we think that the local accuracy constraint is a genuine constraint on rationality, we don't think the global accuracy constraint is.

In terms of generality, there are general constraints that are neither local nor global:

Generalised accuracy constraint: It is not rationally co-tenable to hold your F-beliefs and believe there is an F-belief that is inaccurate.

We can pick out sets of beliefs in various ways. The 'F-beliefs' might be grouped by source (e.g., beliefs based on visual experience), content (e.g., beliefs about Dolly Parton's career), epistemic status (e.g., your rational beliefs), and so on. For the same reasons we think the global accuracy constraint doesn't hold, we think that some generalised accuracy constraints won't hold.

In some cases, the failure of the generalised or global accuracy constraints will be connected to failures of generalised or global evidence constraints:

Local evidential constraint: It is not rationally co-tenable to believe p and believe your belief about p isn't supported by the evidence.

Global evidential constraint: It is not rationally co-tenable to believe what you do and believe there is something you believe that is not supported by the evidence.

Generalised evidential constraint: It is not rationally co-tenable to hold your F-beliefs and believe there is an F-belief that is not supported by the evidence.

We'll consider some of the cases in a moment.

We want to note that there should also be some metacoherence constraints that connect beliefs and credences. Consider the global accuracy constraint, the constraint that says that it's not rationally co-tenable to believe what you do while believing that there is something you believe that's false. Compare this to this constraint governing credences and beliefs: the constraint that says that it's not rationally co-tenable to believe what you do about whales, say, whilst being nearly certain that nearly everything you believe about whales is mistaken. We think that if the *outright* belief puts rational pressure on the thinker to suspend (to the extent that this belief is rational), the *credences* in light of which a thinker rationally expects widespread epistemic failure should also put pressure on the thinker to suspend.

Here is a generalised accuracy constraint that we think can pose trouble for the explanationist:

(*) It is not rationally co-tenable to hold your directly justified beliefs and believe that at least one of your directly justified beliefs is false.

This constraint uses some new jargon. The *directly* justified beliefs are those beliefs that, according to the explanationist, are justified by virtue of being parts of the best explanation of the thinker's evidence. (The indirectly justified ones are only justified by virtue of being explanatory consequences of such beliefs and we'll set those aside.)

We can ask two questions about (*). Is (*) a genuine metacoherence constraint? Does the explanationist view predict that (*) is a genuine metacoherence

constraint? We think that (*) is *not* a plausible metacoherence constraint, but we don't see how the explanationist can take a plausible line on constraints like (*). On the one hand, we think that a thinker can acquire evidence that makes it rational to believe she holds at least one directly justified belief that is false. On the other, think the explanationist either gets this wrong or gets something very similar to (*) wrong. If they end up saying that (*) is a genuine constraint, their view predicts that evidence that shouldn't defeat threatens a large class of beliefs *en masse*.¹⁵ If they somehow dodge this result, they'll be committed to a view that treats evidence that *should* defeat as if it's rationally benign.

Here's an initial case. Here's the first step. Imagine an author has written a large work of non-fiction that contains all and only her directly supported beliefs. At this stage, upon pain of scepticism, it must be possible that this book contains a great many of her beliefs about contingent matters of fact that she might rationally believe without being rationally certain that each of those things is true. Here's the second step. A panel of experts tells her that her book contains *precisely* one error. Here's the third stage. Our author must 'decide' whether to expand her book by adding this to her preface and whether to revise her book by removing some of its content.¹⁶ Remember that the book should contain all and only her directly supported beliefs, so at this stage, we can ask this: can she have direct support for her belief that the book contains an error *and* (continued) direct support for each of the beliefs in her book?

The explanationist view tells us that a belief is directly justified iff it is part of the best explanation of the evidence and that a belief will cease to be directly justified if it ceases to be part of some best explanation of the evidence. Once we know or rationally believe that an explanation contains a falsehood, the beliefs contained in that explanation will cease to be directly justified because no successful explanation contains falsehoods. (Once it's certain that something isn't a successful explanation, it's no longer one of the best and good enough explanations.) The explanationist can either say that this information about the success of the

¹⁵ In what follows, we will assume that the beliefs contained in the aggregate are supported equally and are equally doubtful so that when there is rational pressure to suspend on the aggregate there are not beliefs that constitute discernibly weak elements in the collection that we might dispose of while retaining the rest. For a helpful discussion of this assumption in preface-type cases, see Smith (forthcoming).

¹⁶ This version of Makinson's (1965) case can be found in Littlejohn & Dutant (2020) and Praolini (2019). Ryan (1996) and Smith (2022) provide important discussions of the preface, including ones with the stipulations from above. We think that while Smith's (2016) view is similar to explanationism in some respects, the main difference being that the explanationist isn't so clearly committed to the idea that our evidence always supports a set of propositions that's logically consistent. We provide further arguments against imposing this consistency requirement in Littlejohn & Dutant (forthcoming) where we argue that views that impose this requirement seem to either impose a certainty requirement on rational belief or won't fit with our intuitions about epistemic desirability.

explanations contained in the book (a) makes it the case that not each of the initially directly justified beliefs remains so or (b) each of the initially directly justified beliefs might remain so.

Let's consider the (a) answer first. On this view, once the author learns that the book contained an error and learns that one explanation is unsuccessful, the beliefs contained in one or more of the explanations ceases to be justified. They can either say that some are lost or that each is lost. We'll suppose that each explanation antecedently seemed to be on equal footing. Without any discernible weak link, it seems that the view that says that some remain justified and others do not is too externalist in the sense that the support relation picks winners and losers on grounds that couldn't be discernible to the thinker. (Alternatively, we could think of this selection as arbitrary.) On the other hand, the view that doesn't choose between winners and losers is too sceptical. For reasons we'll sketch below, it seems that learning that one and only one mistake was made shouldn't lead to the widespread abandonment of beliefs that were adequately justified previously.

If this is right, then it seems the explanationist might opt for (b), the view that each of the beliefs that were directly justified remain so. We assume that the explanationist doesn't want to say that this is so because the author isn't justified in believing her work contains an error. It seems dogmatic to retain beliefs and refuse to believe a panel of experts when they tell you that you've made a mistake. If the explanationist says that the author can rationally believe that the book contains a mistake *and* retain belief in each of the parts of the original explanations, they can reject (*). We think that anyone who believes there can be rational but false beliefs should reject (*), but we worry that the explanationist who goes this route will struggle with related constraints like this:

(**) It is not rationally co-tenable to hold your directly justified beliefs and believe that many/most/nearly every one of your directly justified beliefs is false.

We can imagine lining up increasingly error-ridden books similar to the authors where the number of errors is revealed by the panel. Once we arrive at those books where nearly every one of the directly justified belief is believed to be mistaken (where this belief is based on the panel's testimony), we should find a case where the evidence *cannot* provide adequate support for each of the initial beliefs *and* the belief about the prevalence of errors and unsuccessful explanations.

How does the explanationist view make sense of the difference along this continuum (i.e., the difference between (*) and (**))? We can see how in a local case the (apparent) discovery that an explanation fails can defeat the justification to believe the claims contained in the explanation, but we cannot see how this information about the ratio of success to failure could bear on the particular explanations the author initially accepted. Suppose she initially believed something about octopi, something about WWI, something about the clarinet, something about whiskers on kittens, etc. by means of justification-conferring explanatory inferences only to then learn that most of these explanations were unsuccessful. We don't see how this bears on the loveliness of the various explanations for the evidence about octopi in particular. We do think we see, however, that information about the ratio

of success to failure matters to the rationality of this author's beliefs—when the risk of believing falsehoods is too great, it seems obviously irrational to believe even *if* that belief is contained in some seemingly lovely explanation.

It is no mystery why scale would matter on a strength-centred view. There are some risks that we rationally should tolerate and some we should not. Deciding which risk is which is, according to the strength-centred theorist, determined by the probability of undesirable outcomes and the values involved. It seems that the explanationist view offers us no rational basis for distinguishing (*) from (**), so perhaps they'll encourage us to bite the bullet and accept both (putative) constraints.

Is it really so bad to accept (*)? We think so. Our case is a version of Makinson's (1965) preface with a few twists.¹⁷ It's important to our argument that readers agree that the preface differs from the lottery in that it's possible to have a preface-type case in which there's a set of beliefs that's known to be inconsistent and still be comprised entirely of rational beliefs.

One thing to note about the preface-type cases is that some preface-type cases seem to be paradigmatic cases of knowledge. In a preface case in which (say) a subject memorises every entry in a phone book along with the claim that the phone book contains an error, the subject seems to acquire quite a lot of knowledge about phone numbers and knowledge of the fallibility of the source. In the admittedly odd case where an expert testifies the book contains one error out of n claims, we think, in principle, someone could come to believe each of the n claims and the true ones (i.e., $n-1$ claims) could be known. In the absence of discernible differences between the n -cases, it is very tempting to say that the beliefs will not differ in terms of their rational status. Paradigmatic cases of knowledge are plausible cases of rational belief. Cases in which it's nearly certain that someone will come to know are also plausible cases of rational belief. Denying knowledge in these cases leads us to something close to scepticism.

We can put further pressure on views that accept (*). When we accept (*) and try to avoid the apparent sceptical consequences of (*), we end up with very odd views about comparative epistemic preferability about sources or about responses. Contrast (*) with this:

(***) It is not rationally co-tenable to hold your directly justified beliefs and harbour small doubts about the accuracy of each of these beliefs.

Here's a toy case. Suppose we say that if a thinker has a credence of .95 or greater in p , this thinker's doubts about p are small. Now consider two questions about epistemic preferability, a question about preferable situations/sources and a question

¹⁷ We think that while Ryan (1991) is right that it might be rare that the book consists entirely of well-founded beliefs, we should be able to focus on the hypothetical case where this condition is met. We also think that it's a distraction to formulate the case in such a way that the evidence for believing that the relevant set of beliefs or claims contains a falsehood is statistical. The move to expert testimony should remove the temptation to say that our attitude towards the proposition that the relevant set contains an error should be something other than outright belief.

about preferable responses. In our original case (Author 1), the author believed n claims (including the claim that $n-1$ claims were true). In our variant case (Author 2), the author is not told by a panel that she's made any mistakes. Instead, she believes n claims and is .95 confident in each. Let's suppose the number of claims here is 100. Notice that in Author 2, the expected number of errors greatly exceeds the number of expected number of errors in Author 1. In Author 1, the expected number of errors matches the believed number of errors and actual number of errors: 1. In Author 2, the expected number of errors is 5. If we accept (*) and reject (**), we're saying that suspension is preferable to belief in (*) but not in (**). Why would we prefer believing in the case where the expected number of errors and failed explanations is *greater* than in the case where suspension is preferable to belief? This preference strikes us as odd.

Ask yourself which source you would prefer. We cannot think of any reason why anyone would prefer the second author's work as a source to the first. It's true that nobody is yet *convinced* that the second source contains an error, but it's hard to believe that this difference between conviction and expectation matters much. If we told you (and you believed us) that out of 100 fireworks, one will be a dud and then asked you whether you would trade this for a box of 100 where the expected number of duds is 5, it's hard to believe that if you wanted to get a bang for your buck, you wouldn't gladly trade the former for the latter.¹⁸

If the explanationist responds by saying that we should accept (**), we feel that the game is up. If we cannot rationally believe when we harbour only these small doubts about accuracy, we're going to be forced to move towards the view on which none of our beliefs are rational unless it's completely certain that they are true. But this view, we think, is functionally indistinguishable from scepticism and impossible to reconcile with the generally optimistic outlook epistemologists take towards the prospect of acquiring knowledge from fallible sources like phone books, Wikipedia entries, and friends who are not endowed with unnatural abilities.

3. Strength and knowledge

We seem to be pulled in different directions by our intuitions. When we try to make sense of the intuitions that support Seeing, Hearing, and Playing, we might be tempted to embrace explanationism and reject High. When we try to make sense of intuitions like these, a strength-centred view might seem attractive:

Authoring: if an author carefully researches her large book and bases each claim on the right kind of evidence, she can rationally believe each of the claims in the book, including the claim that the book contains an error.

¹⁸ Assuming, that is, that you don't need for some reason to get 100 explosions. If you value each explosion equally and are disappointed by each dud equally, it's clear to us that when presented with a choice between a container of n items where it's certain that $n-1$ are desirable and a container where the expected number of undesirable items greatly exceeds one, you should prefer the former.

Authoring-: Even if the author carefully researches her book and bases each claim on the right kind of evidence, it can be irrational for her to retain belief if the expected number of errors is too great.

It's tempting to appeal to a strength-centred theory to explain Authoring and Authoring- since the most obvious difference between the cases where it is rational for the author to believe and the case in which it is not is precisely that the strength of support decreases as the number of expected errors increases.

We think that we can make the most progress in making sense of these intuitions by introducing a strength-centred view that explains Low and explains why we should reject High. Our view will be strength-centred by virtue of the fact that it incorporates expectationism. It differs from the most familiar strength-centred theories in that it doesn't combine expectationism with veritist assumptions. The key, we think, to making sense of these puzzling intuitions is to retain expectationism and combine that view with a better view of what's truly epistemically desirable.

Many theories of rational belief fail to predict or explain the intuitive difference between lottery-type cases and preface-type cases, but if we're right about Playing and Authoring, this is something that a theory of rational belief should be expected to do. Foley observed that there is at least one key difference between the cases:

To be sure, 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 (2009: 44).

Unfortunately, he immediately added the remark that, "This difference, however, is to be explained by the prerequisites of knowledge, not those of rational belief" (2009: 44). We don't see why this difference should be irrelevant to the theory of rational belief if rational believers desire to *know* the truth and have an aversion to believing without knowing the truth.

Suppose that what's desirable from the epistemic point is acquiring knowledge and that it's undesirable to believe without thereby acquiring knowledge so that we replace veritism with this theory of epistemic value:

Gnosticism: It is most epistemically desirable (objectively speaking) to acquire knowledge and the least desirable outcome is believing without knowing. Suspension is preferable to belief if that belief fails to constitute knowledge.

What happens if we opt for a strength-centred view that combines expectationism with gnosticism instead of veritism? We end up with this view:

Gnostic expectabilism: It is rational for a thinker to believe *p* iff (and because) it is rational for this thinker to be

sufficiently confident that by believing p , she will know p .¹⁹

On the supposition that it would be worse to believe what's not known than it would be to fail to believe in a situation in which a belief would constitute knowledge, gnostic expectabilism tells us that it's rational to believe p outright only if it's more likely than not that by so believing one will come to know.

When it comes to Seeing, Hearing, and Playing, gnostic expectabilism and explanationism deliver the same verdicts. In the lottery case, it is very likely that the lottery proposition is true, but it is *certain* that the belief will not constitute knowledge. Given the certainty, it is surely not true that it's more likely than not that by believing a thinker will acquire knowledge and so suspension does better than belief in terms of expected epistemic value. In mundane cases, however, we normally assume that the cases in which it's rational to believe on the basis of testimony, memory, or observation are cases in which it's not likely that we'll fail to acquire knowledge. When it comes to these intuitions, it seems that explanationism and gnostic expectabilism are both doing fine.

It should be noted that we now have on the table a strength-centred view (admittedly, an unorthodox one) that explains Low and explains why we should reject High. If it's rational to believe only if it's more likely than not that you know, it's rational to believe only if it's more likely than not that your belief is correct. (The probability of knowing will generally be exceeded by the probability of the truth of the target proposition.) We know that High fails, on this view, because however high the probability of the target proposition, the probability that this proposition could be known could be 0. At the very least, then, we can see that the lottery argument against strength-centred views only works against a subset of such views. It fails as an argument against all such views.

When it comes to Authoring and Authoring-, we think gnostic expectabilism delivers just the verdicts we want.²⁰ As the expected number of errors

¹⁹ In Littlejohn & Dutant (2021), we use this view to give a unified treatment of defeaters. On this view, defeaters are defeaters because they are 'indicators of ignorance' (i.e., evidence that lowers the probability that our beliefs meet the conditions necessary for knowledge). We provide further arguments for this approach in Littlejohn & Dutant (forthcoming).

²⁰ We should note that some might be concerned that unsophisticated agents (e.g., animals, children) might not have the cognitive tools to track their degrees of confidence. Is this a problem for our view? We think that if it is, it's a problem that arises for any expectationist view, including the Lockean view. The main difference between their proposal and ours is really about the desirable properties of belief. Of course, someone might think that this is to their advantage because (a) rationality is not a constitutive part of the property they take to be desirable (i.e., accuracy) and (b) rationality is a constitutive part of the property we take to be desirable (i.e., knowledge). We have a quick response to this. We deny that rationality is necessary for knowledge. Actually, the case of unsophisticated agents is helpful here. If I tell you Agnes knows that it's time for dinner, you might *think* she rationally believes

increases, our confidence that the beliefs that correspond to the claims in the book constitute knowledge will decrease. Even if it's a possibility that such beliefs 'turn out' to be knowledge, belief might do worse than suspension in terms of expected epistemic value and when that happens, we can see why the scale matters and why we should take the view that there's a continuum of cases here. We should not take the view that a handful of errors is incompatible with rationality since it's compatible with having most of the beliefs corresponding to claims in the book turn out to be knowledge. We should not take the view that widespread inaccuracy is rationally tolerable. We don't see how a view that dispenses with expectationism can explain the full range of cases here, so we think that our expectationist view does a better job with preface-type cases than the explanationist view can.

Gnostic expectabilism and explanationism take explanatory considerations to matter to rationality. We want to note some of the differences in the ways that we do that. It is supposed to be a platitude about knowledge that knowledge differs from mere true belief in that knowledge requires that it's not a mere coincidence that the belief in question is true.²¹ If this is right, some explanatory connection between the facts that our beliefs concern and our beliefs must hold when those beliefs constitute knowledge. On our view, the presence of such an explanatory tie is a necessary condition for the realisation of the fundamental epistemic good that rational believers hope to acquire (i.e., knowledge). Its absence, in turn, is a sufficient condition for the realisation of the undesirable epistemic outcome that rational believers hope to avoid, believing without knowing.

Our view does not imply that rational belief requires that the explanatory connection between belief and fact obtains, only that the evidence makes it sufficiently likely that it does. The explanatory connections between belief and fact matter to rationality, we think, but they only matter *indirectly*. Thus, we can have large sets of beliefs where each belief is rationally held where there are some doubts that each belief in that set is connected to the facts in such a way that it isn't a coincidence that they are correct. Our view also might explain why inference to best explanation is itself a way of acquiring rational beliefs. If it's true that this form of inference is a reliable way of acquiring knowledge, it is not surprising that it is a way of forming rational beliefs. Doubts about whether *all* beliefs are made rational by virtue of some explanatory inference do not cast doubt on our view since our view explains the rational force of explanatory inference in terms of the expectation that such inferences will expand our knowledge. Since we don't think it's wise to reject

this or that she believes what she should. If I tell you that she's an infant or a dog, you might agree with the knowledge attribution but feel funny about the attribution of rational belief. Once we think of knowledge as a kind of non-normative relation between an animal and a fact (Hyman (1999)), it's eligible to play the role we're assigning to it, that of the desirable property that bears an indirect relation to rationality. For defence of this non-normative notion of knowledge, see Kornblith (2002) and Sylvan (2018). Thanks to Arturs Logins for raising some of these issues.

²¹ For helpful discussions of the connection between knowledge and explanation, see Jenkins (2006) and Nelkin (2000).

expectationism, we don't think that the explanationist is right to put explanation into their theory of how we ought to rationally pursue our epistemic ends if we want to acquire true belief or knowledge (and avoid false belief or ignorance). We also don't see why someone would want to treat loveliness of explanation as a kind of epistemic prize that supplants the role that we assign to knowledge. We think it makes more sense to think that things turn out well when we acquire knowledge and that things turning out well require a lovely explanation only in those cases where the explanatory inference was our means for coming to know. So, while we find some room for explanatory considerations in our theory of rational belief, it differs from the role that the explanationist assigns to it. We think this helps us see why we'd approach the lottery cases in similar ways and why our approach might have important virtues when it comes to metacoherence constraints and preface-type cases of the kind discussed above.

4. Prizes and Pursuits

We know that some readers might be sceptical of the gnostic value theory. We cannot offer a full defence of that here. We wanted to note, however, a few things in its support.

First, the theory fits with things that epistemologists often say in the course of trying to explain the value of truth. They often describe the desire for truth in terms of a desire for *having* the truth or a desire for being in touch with reality. This desire to be in touch with reality or having the truth is then identified with the desire for having true beliefs, but it's not at all clear that these come to the same thing.

Think, for example, of Nozick's (1974) experience machine. In the machine, we think that a subject is completely cut off from reality and, as Lynch (2004) puts it, what we desire in desiring the truth is a kind of contact with reality that's missing from the life we'd live if trapped in the machine. Being in the machine seems to be incompatible with being in touch with external realities. What Nozick seems to be right about is that we think that being regularly in contact with certain external realities is necessary for attaining goods that we prize quite highly and, arguably, take to be necessary for living a life worth living. In the machine, we'll find no friendship or love and our projects come to nothing.

Here's a conjecture. If we're focusing on just beliefs about the world beyond appearances, the fundamental epistemic good that we seek is unattainable in the experience machine. We learn from Gettier that being trapped in the machine and having experiences produced by a machine that doesn't take account of what's happening in reality might trigger true beliefs (e.g., it might be a coincidence that the machine produces the experiences that convince you that your political party won the recent election right when your political party happens to win the recent election). When you come to believe that your party won, you believe the truth. Are you thereby 'in touch' with the events taking place in the world that you correctly believe are happening because of this bizarre coincidence? No. Is this really what you want, this kind of mere match between fact and belief that could amount to a mere coincidence in which you are detached from reality? Maybe not. Maybe only

knowledge gives us the connection we desire. That intuition favours gnosticism (Littlejohn, 2013).

Our impression is that many philosophers think that something else must be able to give us this connection. For someone like McDowell (1995), we need some kind of contact with reality prior to and independent from knowledge in order to acquire knowledge. This is because, he thinks, knowledge requires the possession of reasons that guarantee that we can know.²² To be sure, they might say, mere true belief doesn't give us the connection we seek. The experience machine tells us that. It's striking, however, that in the experience machine we're also not in *perceptual* contact with our surroundings. Won't perceptual contact give us the connectedness to the world we seek?

We think not. Nozick's thought experiment suggests to some that there are *relational* goods, goods that we can enjoy only if we bear the right relation to things outside of us where this relation has some psychological or mental dimension. We don't think mere perceptual awareness gives us this connection for two reasons. First, such perceptual relations might not reveal the meaning or significance of these events for an individual. If it's perceptual contact rather than knowledge that explains why it matters that you are connected to some event of tremendous personal importance (e.g., a wedding, a funeral), remember that the perceptual relations you bear to these events will be similar to the ones that children and animals bear to these events. It matters how we conceive of the events we have before the mind and we think that perceptual contact places too few constraints on how events are presented to do the requisite work.²³

²² In his recent work, Schroeder (2021) has also suggested we need contact with reasons provided by something other than knowledge because knowledge itself should be understood in terms of reasons. On his view, knowing is a matter of believing for sufficiently good objective and subjective reason. His view is less epistemically demanding than McDowell's in that Schroeder's view (rightly, we think) allows for the possibility of knowing things on the basis of reasons that are fallible. On McDowell's view, reasons are adequate only when a complete description of them entails that we're in a position to know. That's stronger than infallibilism and we think infallibilism is too strong already. What Schroeder and McDowell urge us to recognise is that we must have a way of possessing reasons that doesn't involve knowledge if we're going to explain how beliefs 'turn into' knowledge by being based on reasons. That conditional is impeccable, but we don't think that all knowledge must be based on propositionally specified reasons.

²³ Obviously, this issue requires further discussion about the connection between perception and cognition. Some authors, particularly Silva (2020, n.d.), have argued that we can be aware that something is so even if we do not know that it's so. Presumably, being aware that someone is being buried or wed requires conceiving of the relevant events under certain descriptions. We're somewhat sceptical of the claim that a thinker can be aware that *p* without knowing that *p* (and the same scepticism extends to the claim that we can perceive or remember that *p* without knowing that *p*). Those who insist that we *can* have propositional awareness without

Second, perceptual relations are *neutral* on what's happening outside of us. If someone is *upset* by something, seeing something upsetting isn't sufficient, not if you're agnostic about whether this event is happening. The realisation of some relational goods requires *conviction* that is contained in belief but not contained in perceptual consciousness. Knowledge gives us conviction and connection. Perception, at best, gives us a kind of connection.

Things happening in the world which we're completely detached from might not harm or benefit us. This thought, we think, helps to explain why people find the prospect of life in the machine unpalatable. If knowledge is necessary for being related to these events so that they might be good or bad *for us*, we think that we'll make some progress towards explaining the value of knowledge and the connection between knowledge and belief. Belief's role, we think, includes that of putting us in contact with reality, so beliefs that fail to do that are, to that extent, undesirable. Our hypothesis is that beliefs that fail to constitute knowledge are undesirable for this reason. Once we're clear that there's a notion of objective epistemic desirability on which desirability involves more than a 'mere match' between believed propositions and the world, it starts to make sense that more than mere high probability is necessary for rational belief. It also seems that the project of finding new and exotic connections between belief and truth is somewhat misguided if that doesn't ultimately appeal to the properties that beliefs need to be knowledge.

We think that epistemically loaded prizes (i.e., prizes that can only be completely and accurately described by reference to some thinker's epistemic state) have been largely overlooked in the literature on epistemic rationality. We also find this neglected in the literature on practical rationality. Once we start to think that *knowing* things about your surroundings or people you surround yourself with is necessary for realising certain goods (e.g., being in a loving relationship, sharing meaningful projects, developing true friendship), it makes sense that in some choice situations, information about what we might know or not know matters to choice.²⁴ Knowledge might matter to rational choice if, say, what we know determines in whole or part what evidence we have. Rational agents don't just want to change the worlds in certain ways without being informed of such changes. We think rational agents sometimes desire for the world to change in certain ways whilst enjoying awareness of these changes. If our preferences are sensitive to such epistemically loaded prizes, knowledge can matter to rational choice because of the role it plays in describing prizes and needn't play any role in the theory of how rational agents pursue the things they value.

propositional knowledge often suggest that we don't need *belief* for propositional awareness. That fits with the role they assign to propositional awareness, but it points to a problem we take to be fatal for these proposals. We cannot be guided by facts or apparent facts in the relevant ways if we're agnostic or neutral.

²⁴ For arguments that knowledge is an essential part of certain emotional reactions and reactive attitudes, see Logins (forthcoming) and Unger (1975).

5. Conclusion

We have proposed a theory of rational (full) belief that seemingly does the impossible. It reconciles our intuitions about lottery cases with a strength-centred theory of rational support. The key to reconciling this attractive approach to thinking about the sufficiency of sufficient support with our intuitions about lottery cases is to think about what a thinker's evidence has to 'say' to her to convince her to take a stand on some matter. We don't think it's enough to produce conviction in addition to high credence that the evidence says, 'Hey, p is probably true!' If the evidence says *both* 'Hey, it's very likely that p ' and 'You'll never know whether p ', we don't think that a rational thinker should therein become convinced that p is true. If the evidence, however, says, 'Hey, it's nearly certain that you'll know the answer to the question whether p if you believe p ', we think a rational thinker should settle question whether p affirmatively. That's because we think such thinkers rationally aspire to know whether p and so desire to know and are averse to believing what's not known.

What the evidence must do to convince a rational thinker to commit is say that the risk of believing without knowing is sufficiently small. The evidence is sufficient for rational belief when it provides sufficiently strong support for the hypothesis that the thinker will come to know by believing. This view combines a not wildly unpopular value theory and a wildly popular way of thinking about rationally responding to situations that involve an element of risk and delivers a non-standard theory of rational belief that seems to vindicate intuitions old and new that cause trouble for alternative approaches. While this might be sufficient to convince readers that our proposal is correct, we hope it will convince them to give our proposal serious consideration and think of whether there are alternative accounts that better make sense of the data.²⁵

References

- Appley, Bryan C., and Gregory Stoutenburg. (2017). Two new objections to explanationism. *Synthese*, 194(8), 3069–3084. <https://doi.org/10.1007/s11229-016-1093-1>
- Bird, Alexander. (2007). Justified Judging. *Philosophy and Phenomenological Research*, 74(1), 81–110. <https://doi.org/10.1111/j.1933-1592.2007.00004.x>
- Climenhaga, Nevin. (2017). Inference to the Best Explanation Made Incoherent. *The Journal of Philosophy*, 114(5), 251–273. <https://doi.org/10.5840/jphil2017114519>
- Conee, Earl. (2004). Externalism, Internalism, and Skepticism. *Philosophical Issues*, 14, 78–90.
- Conee, Earl Brink, and Richard Feldman. (2004). *Evidentialism: Essays in Epistemology* (Vol. 48, 15–34). Oxford University Press.

²⁵ We want to thank Artūrs Logins, Jacques-Henri Vollet, and an anonymous referee for very helpful comments on this paper. We also want to thank Kevin McCain and Martin Smith for discussing the issues in this paper.

- Dellsén, Finnur. (2021). Explanatory Consolidation: From 'Best' to 'Good Enough'. *Philosophy and Phenomenological Research*, 103(1), 157–177. <https://doi.org/10.1111/phpr.12706>
- Easwaran, Kenny. (2016). Dr. Truthlove Or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Noûs*, 50(4), 816–853. <https://doi.org/10.1111/nous.12099>
- Evnine, Simon J. (1999). Believing Conjunctions. *Synthese*, 118(2), 201–227.
- Foley, Richard. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis. In Franz Huber & Christoph Schmidt-Petri (Eds.), *Degrees of Belief* (37–47). Springer.
- Gardiner, Georgi. (2019). The Reasonable and the Relevant: Legal Standards of Proof. *Philosophy and Public Affairs*, 47(3), 288–318. <https://doi.org/10.1111/papa.12149>
- Harman, Gilbert. (1968). Knowledge, Inference, and Explanation. *American Philosophical Quarterly*, 5(3), 164–173.
- Henderson, Leah. (2014). Bayesianism and Inference to the Best Explanation. *The British Journal for the Philosophy of Science*, 65(4), 687–715.
- Hirvelä, Jaako. (2022). Justification and the knowledge-connection. *Philosophical Studies*, 179(6), 1973–1995. <https://doi.org/10.1007/s11098-021-01741-x>
- Hyman, John. (1999). How knowledge works. *Philosophical Quarterly*, 49(197), 433–451.
- Ichikawa, Jonathan Jenkins. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, 44(2), 184–206.
- Jackson, Elizabeth. (2019). Belief and Credence: Why the Attitude-Type Matters. *Philosophical Studies*, 176(9), 2477–2496.
- Jenkins, Carrie. (2006). Knowledge and Explanation. *Canadian Journal of Philosophy*, 36(2), 137–164. <https://doi.org/10.1353/cjp.2006.0009>
- Kelp, Christoph. (2014). Two for the Knowledge Goal of Inquiry. *American Philosophical Quarterly*, 51(3), 227–32.
- Kornblith, Hilary. (2002). *Knowledge and its Place in Nature* (Vol. 115, 246–251). Oxford University Press.
- Lehrer, Keith. (1970). Justification, Explanation, and Induction. In Marshall Swain (Ed.), *Induction, Acceptance and Rational Belief* (100–133). https://doi.org/10.1007/978-94-010-3390-9_6
- Leitgeb, Hannes. (2014). The Stability Theory of Belief. *Philosophical Review*, 123(2), 131–171.
- Lipton, Peter. (2001). Is Explanation a Guide to Inference? A Reply to Wesley C. Salmon. In Giora Hon & Sam S. Rakover (Eds.), *Explanation: Theoretical Approaches and Applications* (93–120). https://doi.org/10.1007/978-94-015-9731-9_4
- Littlejohn, Clayton. (2013). The Russellian Retreat. *Proceedings of the Aristotelian Society*, 113, 293–320.
- Littlejohn, Clayton. (2015). Who Cares What You Accurately Believe? *Philosophical Perspectives*, 29(1), 217–248.
- Littlejohn, Clayton, and Julien Dutant. (2020). Justification, Knowledge, and Normality. *Philosophical Studies*, 177(6), 1593–1609.

- Littlejohn, Clayton, and Julien Dutant. (2021). Defeaters as Indicators of Ignorance. In Mona Simion & Jessica Brown (Eds.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Littlejohn, Clayton, and Julien Dutant. (forthcoming). What is Rational Belief? *Noûs*.
- Logins, Artūrs. (forthcoming). Persistent burglars and knocks on doors: Causal indispensability of knowing vindicated. *European Journal of Philosophy*, 1–23.
- Lycan, William G. (1988). *Judgement and Justification*. CUP Archive.
- Lynch, Michael P. (2004). *True to Life: Why Truth Matters*. MIT University Press.
- Makinson, David C. (1965). The paradox of the preface. *Analysis*, 25(6), 205.
- McCain, Kevin. (2014). *Evidentialism and Epistemic Justification*. Routledge.
- McCain, Kevin, and Luca Moretti. (2021). *Appearance and Explanation: Phenomenal Explanationism in Epistemology*. Oxford University Press.
- McDowell, John. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, 55(4), 877–93.
- Moss, Sarah. (2018). *Probabilistic Knowledge*. Oxford University Press.
- Nelkin, Dana K. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, 109(3), 373–409.
- Nozick, Robert. (1974). *Anarchy, State, and Utopia* (187–201). Basic Books.
- Owens, David. (2000). *Reason Without Freedom: The Problem of Epistemic Normativity*. Routledge.
- Pollock, John L. (1986). The Paradox of the Preface. *Philosophy of Science*, 53, 246–258.
- Poston, Ted. (2011). Explanationist Plasticity and the Problem of the Criterion. *Philosophical Papers*, 40(3), 395–419. <https://doi.org/10.1080/05568641.2011.634248>
- Praolini, Francesco. (2019). No Justificatory Closure without Truth. *Australasian Journal of Philosophy*, 97(4), 715–726. <https://doi.org/10.1080/00048402.2018.1564059>
- Rosenkranz, Sven. (2021). *Justification as Ignorance: An Essay in Epistemology*. <https://doi.org/10.1093/oso/9780198865636.001.0001>
- Russell, Bertrand. (1912). *The Problems of Philosophy*. Oxford University Press.
- Ryan, Sharon. (1991). The preface paradox. *Philosophical Studies*, 64(3), 293–307.
- Ryan, Sharon. (1996). The epistemic virtues of consistency. *Synthese*, 109(2), 121–141. <https://doi.org/10.1007/BF00413765>
- Schroeder, Mark. (2021). *Reasons First*. Oxford University Press.
- Shah, Nishi. (2006). A new argument for evidentialism. *Philosophical Quarterly*, 56(225), 481–498.
- Silva, Paul. (2020). Possessing Reasons: Why the Awareness-First Approach is Better Than the Knowledge-First Approach. *Synthese*, 199(1–2), 2925–2947. <https://doi.org/10.1007/s11229-020-02916-5>
- Silva, Paul. (n.d.). Basic knowledge and the normativity of knowledge: The awareness-first solution. *Philosophy and Phenomenological Research*, n/a(n/a). <https://doi.org/10.1111/phpr.12754>
- Smith, Martin. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press UK.
- Smith, Martin. (2022). The Hardest Paradox for Closure. *Erkenntnis*, 87(4), 2003–2028.

- Sutton, Jonathan. (2005). Stick to What You Know. *Noûs*, 39(3), 359–396. <https://doi.org/10.1111/j.0029-4624.2005.00506.x>
- Sylvan, Kurt. (2018). Knowledge as a Non-Normative Relation. *Philosophy and Phenomenological Research*, 97(1), 190–222.
- Unger, Peter K. (1975). *Ignorance: A Case for Scepticism* (369). Oxford University Press.
- Williamson, Timothy. (2000). *Knowledge and its Limits*. Oxford University Press.
- Worsnip, Alex. (2016). Belief, Credence, and the Preface Paradox. *Australasian Journal of Philosophy*, 94(3), 549–562. <https://doi.org/10.1080/00048402.2015.1084343>