

# What is rational belief?

Julien Dutant<sup>1</sup>  | Clayton Littlejohn<sup>2</sup>

<sup>1</sup>King's College London

<sup>2</sup>Dianoia Institute of Philosophy

## Correspondence

Clayton Littlejohn, Dianoia Institute of Philosophy.

Email: [cmlittlejohn@gmail.com](mailto:cmlittlejohn@gmail.com)

## Abstract

A theory of rational belief should get the cases right. It should also reach its verdicts using the right theoretical assumptions. Leading theories seem to predict the wrong things. With only one exception, they don't accommodate principles that we should use to explain these verdicts. We offer a theory of rational belief that combines an attractive picture of epistemic desirability with plausible principles connecting desirability to rationality.

## 1 | INTRODUCTION

A theory of rational belief should get the cases right<sup>1</sup>. It should predict intuitively compelling verdicts or help us see things in a new light. It should also incorporate the principles we should use to explain these verdicts. While extensional adequacy matters, intuitive verdicts are not our only guide. A theory of rational belief should incorporate sound principles so that it might say the right things for the right reasons.

Which theories do this? Unfortunately, we don't think that any do. Leading theories predict too little or predict the wrong things. No extant theory gets the cases right. We find only one theory that accommodates attractive principles. Our most promising theories generally say the wrong things for the wrong reasons.

We've stumbled upon a view that we think gets the cases right and can give the right explanation for these verdicts. Readers might think that our theory will eventually meet the business end of a counterexample, but we will show that it is distinctive in that it at least gets the data points discussed below.

We tried to identify an explanatorily adequate theory of rational belief by looking at the problems that arise for three of the most promising proposals in the literature:

**Lockean:** It is rational for you to believe  $p$  iff it is rational for you to be sufficiently confident that you'd believe  $p$  accurately (Dorst, 2019; Easwaran, 2016; Foley, 2009; Sturgeon, 2008).<sup>2</sup>

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2023 The Authors. *Noûs* published by Wiley Periodicals LLC.

**Normality:** It is rational for you to believe  $p$  iff  $p$  is true in the most normal worlds compatible with your evidence (Smith, 2016, 2018).

**Knowledge Duplication:** It is rational for you to believe  $p$  iff you're sufficiently similar on the inside to a thinker who knows  $p$  (Bird, 2007; Ichikawa, 2014; Miracchi, 2015; Reynolds, 2013; Rosenkranz, 2021; Simion et al., 2016; Smithies, 2012).<sup>3</sup>

We focus on these three because each captures an important insight that we should all want our theories to account for. Our theory emerges by considering how we might overcome the challenges that these views face while retaining the insights that made these views attractive. We cannot survey all the alternatives in the literature, but we think that it should be clear by the paper's conclusion how we'd argue against theories we don't discuss here.<sup>4</sup>

The structure of this paper is straightforward. We present objections to the Lockean view, Normality, and Knowledge Duplication in that order. As we proceed, we diagnose the difficulties that arise for these views. We propose our view in §3.

## 2 | LOTTERY LOSSES

Consider:

### Prisoners

One hundred prisoners were exercising in the yard when all but one of them put into action a plan to attack the guard. Our defendant was in the yard when the attack took place, but we have no further information available to determine whether the defendant participated or not (Nesson, 1979).

Here's a data point: it's not reasonable to blame or punish the defendant for joining in the attack (Buchak, 2014; Enoch & Spectre, 2021; Gardiner, 2019; Littlejohn, 2020; Moss, 2018b; Smith, 2018; Thomson, 1986). Focus on attitudinal blame. If we were convinced that the defendant had freely participated in this attack, we would blame them for it.<sup>5</sup> Given the information we have, however, blame isn't reasonable. It might be reasonable if we learned more, but not now. Suspicion might be reasonable. Fear or distrust might be reasonable. Blame seems to require being convinced that the agent has done something bad. The guilty know what they've done and should feel guilt. The doing is the right kind of thing to blame someone for. Thus, it's reasonable to think that it's unreasonable to blame because we don't have the information that would warrant attitudinal blame.

The problem with this information, presumably, isn't that the support it offers isn't sufficiently strong. Replace statistical evidence with an eyewitness and our intuitions about the reasonableness of blame can change even if we should be less confident than in the original case (Harman, 1968; Nelkin, 2000; Smith, 2016). This means there's a problem with the idea that the sufficiency of sufficient evidence for attitudinal blame and the belief states that rationalise it is understood in terms of strength. Bearing this in mind, we have our first argument against the Lockean view:

*Against the Lockean View (I)*

P1. According to the Lockean view, if it's rational to be nearly certain that the defendant is guilty, it's rational to believe outright that she's guilty.

P2. In Prisoners, however, it's not rational to believe this outright even though it's rational to be nearly certain.

C. Thus, the Lockean view is mistaken.

The naked statistical evidence we have in Prisoners isn't enough for rational belief (**Statistical Insufficiency**). It might warrant a high degree of confidence, but nothing more. That's the problem.

We can use Prisoners to formulate a second argument against the Lockean view:

*Against the Lockean View (II)*

P1. According to the Lockean view, it's rational to believe the defendant to be guilty.

P2. It's not rational to believe this when you're in a position to know that you couldn't know whether the defendant was guilty.

C. Thus, the Lockean view is mistaken.

On the Lockean view, the proposition *the defendant is guilty* is sufficiently supported. So is proposition *you couldn't know whether the defendant is guilty* and its conjunction with *the defendant is guilty*. This kind of naked statistical evidence seems to warrant attributions of ignorance that aren't compatible with rational outright belief (**Statistical Ignorance**).<sup>6</sup> You should suspend when you can know you couldn't know whether something is so (Miracchi, 2019; Sosa, 2021). The Lockean view doesn't predict this connection between acknowledged ignorance and suspension.

The intuitions that underwrite the first argument figure prominently in arguments for Normality and Knowledge Duplication. Normality makes quick work of cases like Prisoners. When we have only this statistical information to rely on, the situations in which we're wrong about the defendant aren't less normal than the ones in which we'd judge correctly that the defendant is guilty. Thus, we lack normic support for believing she's guilty. Normality might also account for Statistical Ignorance. It might be that situations in which we rationally believe without knowing are much less normal than the situations in which we rationally believe and know (Goodman & Salow, forthcoming). If so, it might be plausibly said that we cannot have normic support to believe both that *p* is true *and* to believe that we couldn't know whether *p*.<sup>7</sup> If we help ourselves to this not implausible assumption about knowledge, Normality accounts for Statistical Ignorance and Statistical Insufficiency.

Knowledge Duplication tells us that the impossibility of knowing in lottery situations is a reason to think that the beliefs in question cannot be rationally held. As Bird puts it, "knowledge is the standard against which beliefs are compared to see whether they count as justified" (2007: 84). Statistical Ignorance, according to this way of thinking, explains Statistical Insufficiency.<sup>8</sup>

One might worry that we're overlooking a reply on behalf of the Lockean view. The Lockean view tells us (roughly) that rational beliefs are rational because there's a level of support that crosses a threshold. In formulating our objections, we've had to assume that the evidence (a) war-

rants a high degree of confidence and (b) the degree of confidence it warrants is sufficiently high. Moral encroachment complicates things (Basu, 2019; Moss, 2018a). We assumed, perhaps naively, that we should be nearly certain the defendant is guilty in Prisoners. Maybe we shouldn't be nearly certain the defendant is guilty (Gao, 2019; Moss, 2018b). Maybe the degree of certainty warranted is high but only high enough in cases where the practical or moral stakes aren't so significant (Basu & Schroeder, 2019). The Lockean view is off the hook if the Lockeans can deny that we should be nearly certain or have a variable threshold that's sensitive to moral factors (Locke, 2014).<sup>9</sup>

Our response is two-fold. We want a theory of rational belief that's compatible with moral or pragmatic encroachment. Still, we don't appeal to encroachment to explain our verdicts about lottery-type cases. We can strip out the morally salient features of the case and elicit the same intuitions. Consider:

### Rabbits

Four hundred rabbits were held in a large cage. All but one had spots. While hopping around, one of the bunnies toppled the food bowl. We have no further information about which of these adorable bunnies toppled the bowl (Schoeman, 1987).

Was the bowl toppled by a bunny with spots? Probably. We're not convinced, however, that the bunny without spots did not topple the bowl. Naked statistical evidence isn't sufficient to warrant being convinced that the causally responsible bunny had spots, but we don't need pragmatic or moral encroachment to explain this. There are no practical stakes and bunnies aren't people to whom anything epistemic is owed.

The second part of our response is that it's not clear how the Lockeans can use moral encroachment to explain Statistical Ignorance. We think that if it's rational to be convinced that you couldn't know whether  $p$  given the information you might have, suspension is rationally required. We don't see how the Lockeans can appeal to pragmatic factors to explain why the recognition that we couldn't know whether  $p$  requires suspension.

## 2.1 | Lottery Lessons

What lessons should be drawn from the lottery? Smith (2010: 11) thinks we should recognise that this principle is false:

**Comparative Accuracy:** If it's rational for you to believe  $q$  and it is more probable on your evidence that  $p$  is true than  $q$  is, it's rational for you to believe  $p$ .

Along with Comparative Accuracy, he rejects the idea that rational beliefs are rational whenever the risk of believing in error is sufficiently small.

We agree. We think this is an important insight that our theories should capture. Before we discuss the alternatives to the Lockean view, however, we should revisit an argument for Comparative Accuracy (Dorst, 2019; Easwaran, 2016; Pettigrew, 2016). We think it's important to try to pinpoint where this argument goes wrong.

Comparative Accuracy follows from these two claims:

**Comparative Desirability:** If it's rational to believe  $q$  and believing  $p$  is more likely to be objectively epistemically desirable than believing  $q$ , it's rational to believe  $p$ .<sup>10</sup>

**Veritism:** A belief is objectively epistemically desirable iff it is true. It is undesirable to believe  $p$  otherwise. Suspension and lack of belief fall somewhere in between (Goldman, 1999; Pettigrew, 2015; Whiting, 2013).

If we reject Comparative Accuracy, we must reject one of these claims. Smith seems to reject Comparative Desirability. This suggests he might accept this diagnosis of the lottery case:

Even if we think of true beliefs as desirable and false beliefs as undesirable, it doesn't follow that rational beliefs are rational because they maximise expected objective epistemic desirability. Rationality requires a different connection to truth, one encoded in Normality.

Smith wouldn't be alone in rejecting Comparative Accuracy without rejecting Veritism. Normality isn't the only extant theory that seeks to explain our lottery intuitions without abandoning Veritism.<sup>11</sup> Even if rational beliefs are rational because connected 'properly' to the truth, the connection between rationality and truth might require something different to the small risk of error.

Consider an alternative response to the argument for Comparative Accuracy. We reject Veritism in favour of an alternative understanding of epistemic desirability:

**Gnosticism:** A belief is objectively epistemically desirable iff it constitutes knowledge. It is undesirable to believe  $p$  otherwise. Suspension and lack of belief fall somewhere in between (Gibbons, 2013; Littlejohn, 2013; Smithies, 2012; Williamson, 2000).<sup>12</sup>

On this view, accuracy is necessary but insufficient for objective epistemic desirability. This is a natural view to adopt if you think, as Bird (2007) does, that knowledge is the standard against which belief is measured. Gnosticism might seem attractive if you think that our curiosity isn't satisfied unless we know (Whitcomb, 2010). Belief might differ from credence in terms of its connection to curiosity (e.g., the evidence that shouldn't satisfy our curiosity might be insufficient for outright belief, a state we're in when our curiosity is sated). It might seem attractive if you think that beliefs are supposed to put us 'in touch' with reality and not merely match it. Arguably, beliefs do this when they constitute knowledge but not otherwise (Hyman, 1999). This gives us an alternative diagnosis of where the argument for Comparative Accuracy goes wrong:

Because mere accuracy is not enough for epistemic desirability, we can reject Comparative Accuracy but retain Comparative Desirability. When it's clear that we cannot hope to acquire knowledge, it's not rational to believe even when it's virtually certain that our beliefs will be accurate. We have no reason (yet) to think that beliefs might agree in expected desirability and differ in rational status.

On our diagnosis, it might be misguided to search for some connection between belief and accuracy that holds in preface cases but not in lottery cases is misguided. Even if we found some such relation, doubts about its normative significance might be warranted.

To evaluate these different diagnoses of where the argument for Comparative Accuracy errs, we should examine Normality and Knowledge Duplication more closely.

### 3 | PREFACE PROBLEMS

Consider:

**Quiz:** You (along with some epistemic peers) were given quiz questions on a diverse range of topics. In each instance, you felt very confident that you knew the answer. This confidence wasn't misplaced. Each question concerned a different topic and you happened to be quite knowledgeable in the relevant areas. After reviewing your answers, the quiz master declares that you knew each of the answers except one'.<sup>13</sup>

Here's another data point: it was reasonable to believe your answers in each case *before* you learned that you made a mistake. Here's yet another data point: you can know and reasonably believe you made a mistake by relying on testimony of the kind provided by the quiz master. Here's a quotation that, we think, captures something important about this sort of case:

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

Ignore the last sentence. Focus on what precedes it. Clear cases of knowledge are plausible cases of rational belief. We think, as Foley does, that it's plausible that it's rational to believe each of the known claims in Quiz *and* that each of the true things believed in Quiz would be known. This includes the beliefs that correspond to the initial answers *and* the belief that there's a flaw hidden among the good answers. For completeness, we'd add that the difference between a true belief (e.g., a belief that corresponded to your belief about chickens) and a false one (e.g., the belief about carbon sequestration efforts in the United States) shouldn't matter to rationality if the beliefs are 'similarly supported'. If this is right, we have our argument against Normality:

#### Against Normality

P1. According to Normality, the set of propositions it's rational to believe at any given time must be logically consistent.

P2. In Quiz, however, it's rational to believe a set of propositions known to be inconsistent because it's rational to believe each of the answers to the quiz questions *and* to believe that one of these answers is mistaken.

C. Thus, Normality is mistaken.

Remember that your evidence provides normic support to those propositions true in the most normal worlds compatible with your evidence. Since there's no normal world in which each proposition in an inconsistent set is true, we get (P1). Our ordinary ascriptions of knowledge coupled with a heuristic linking clear cases of knowledge to rational belief supports (P2).

Have the Lockeans extracted their revenge? Perhaps. They get their revenge against some critics. This line of argument, we think is sufficient to show that rational belief doesn't require normic support. It does not undermine Knowledge Duplication. As Foley observed, the beliefs formed initially were promising cases of knowledge and *remain* promising candidates for being known even after the quiz master informs you that you've made one mistake. Suppose the quiz consisted of 200 questions. You probably didn't expect that you would know 199 of the answers. In each instance, you should be *more* confident that you know the answer than you were initially. The confirmation that your answers contained a flaw was very, very good news. In each instance but one, you were the same on the inside as a knower because you're the same on the inside as yourself. In the remaining instance, you'd be sufficiently similar to someone else who had knowledge.<sup>14</sup>

As with our argument against the Lockean view, we've presented readers with a short argument containing a controversial premise. Smith rejects (P2). He thinks the testimony defeats the rational support possessed initially when the beliefs are similarly supported (Smith, 2022). In denying this, he might also deny Foley's observation that we can know quite a lot in preface-type cases even when we know that there's some false belief hiding among the beliefs that initially constituted knowledge, beliefs that he thinks we should now abandon. How can we settle this disagreement? Should we say that the testimony defeats the rational standing of each belief on the assumption that they are similarly supported? We shall argue that we shouldn't.

### 3.1 | Rationality without Consistency

Let's play two rounds of 'Would you rather?' Modify Quiz slightly. You took the quiz along with nine epistemically gifted friends. Let's suppose further that one of two things might happen after the quizmaster reveals the results:

- (a) You will continue to believe each answer knowing that you've answered precisely one question incorrectly.
  
- (b) You will continue to believe each answer having learned that 9 of you submitted answers that contained 10 mistakes each and that one of you submitted a set of answers that was error-free.

Build into this what you must to convince yourself that each correct answer would be something the relevant thinker would know to be true at least prior to the quizmaster's announcement. Would you rather (a) or (b)?

We're supposing that if you learn that 9 of the 10 sets of answers/beliefs contained 10 mistakes the expected number of errors in each set of answers/beliefs is the same. Given this assumption we conjecture that you'd rather (a) than (b). This suggests this fits your ratherings:

**More than One is Worse:** You'd rather retain your beliefs when it's certain that precisely one of your beliefs is false than retain your beliefs when the expected number of errors greatly exceeds one.

This is compatible with Normality. Normality tells us nothing about which belief set would be worse to retain if retained. It also doesn't tell us that it's acceptable to retain either sets of beliefs. It merely tells us that it's not acceptable to believe each of the things believed initially when it's believed acceptably that one of these beliefs is mistaken.

Round two. If you learned that 9 of the submitted answers contained 10 errors and that one set was error free, would you rather (b) continue to believe or (c) suspend judgment so that you no longer continue to believe any of the things you believed initially? We think you'd rather (b) than (c). This suggests this fits your ratherings:

**More than One is Tolerable:** Even when the number of expected errors greatly exceeds one, you'd sometimes rather believe than suspend.

We think it's sensible to endorse More than One is Tolerable. Why? Suppose your answer key looks something like this:

Q1: What's the capital of California? Sacramento.

Q2: In the War of the Roses, which house battled the Yorks? The Lancasters.

...

Q200: Who was the first woman to win the Nobel Prize? Marie Curie.

Given the similarity of support for each belief, we think that the information about the collection tells us that it's rationally mandatory to suspend in each case only if the information provided gives a compelling reason to suspend when these beliefs are considered individually. On our understanding of the case, there's a reason to suspend when the expected number of errors or expected number of 'ignorant' beliefs (i.e., beliefs that fail to constitute knowledge) is 9 only if there's a reason to suspend on whether Marie Curie was the first woman to win the Nobel Prize given that the probability of knowing the answer isn't greater than 191/200. Given that it seems *eminently* reasonable to believe when it's so likely that the answer is known, we think *this* belief can remain in place. That holds for the others, too.

Of course, our choice of numbers wasn't guided by anything important. Make the quiz 20,000 questions. Make the number of expected errors 4. Maybe 4 greatly exceeds 1. Upon pain of scepticism, we'd better be able to believe when the probability that we know is .9998. Since we can increase the number of questions as much as we like, denying More than One is Tolerable forces us to either (i) embrace the view that rational belief requires being virtually certain that we know or (ii) embrace a view on which the expected inaccuracy of the collection defeats the collected beliefs even when it's virtually certain that the beliefs in question have the properties that are objectively epistemically desirable (e.g., that they are true and that they are known). We find (ii) implausible for the reasons familiar from critiques of epistemic consequentialism. When it comes to propositional attitudes like belief, it's information about the attitudes *qua* particular attitude



and not *qua* attitude contained in a collection that matters to their rational status. This, we think, is what the separateness of propositions intuition shows us (Berker, 2013).<sup>15</sup> We find (i) implausible for reasons familiar from the start. We've been looking at views that predict that we're rational in testimony cases or perceptual cases even when it's not certain that we're right or certain that we know. We're trying to explain the asymmetry between these cases and the lottery without appeal to the idea that testimony or perception provides stronger support than the support we could have in any lottery or any size.

Bearing this in mind, we have our argument that rationality tolerates inconsistency:

### Against Consistency

P1. You'd rather (a) continue believing your initial 200 answers knowing that one is mistaken than (b) continue believing your initial 200 answers when the expected number of errors greatly exceeds one (e.g., upon learning that the expected number of errors is 9).

P2. You'd rather (b) continue believing your initial 200 answers when the expected number of errors greatly exceeds one (e.g., upon learning that the expected number of errors is 9) than (c) suspend on these propositions when the expected number of errors greatly exceeds one.

C. So, you'd sometimes rather believe than suspend even when you know your beliefs are inconsistent (e.g., when you retain your original 200 beliefs and believe one of these beliefs is mistaken).

If your ratherings satisfy a transitivity constraint, we can argue from More than One is Worse and More than One is Tolerable to the conclusion that rationality *can* tolerate believing a set of similarly supported propositions that's certain to contain one falsehood (**Precisely one is Tolerable**). Whilst Normality is compatible with both More than One is Worse and More than One is Tolerable, it must deny Precisely One is Tolerable. We'd rather deny Normality than reject More than One is Worse or More than One is Tolerable. With regards to More than One is Worse, we don't see why having evidence that warrants outright belief that some unidentified belief contains a flaw differs in kind from having evidence that warrants a high degree of confidence that our belief set contains many flaws. Even if we assumed (which we don't) that it's worse to believe whilst being certain that we've erred once on our quiz than it is to believe whilst having one expected error, this acknowledgment that rational toxins are more toxic when believed present as opposed to when feared to be present doesn't undermine our argument.<sup>16</sup>

## 3.2 | Recap

We agree with Smith's (2016) insightful observation that lottery cases are a counterexample to Comparative Accuracy. We worry about proposed alternatives to conflict with Comparative Desirability. The lesson we take from the intuitive difference between lottery and preface cases is *not* that we need to go to the lab and look for increasingly complicated stories about how rationality relates to (mere) accuracy but should instead consider alternatives to Veritism. Gnosticism

seems a natural alternative given that it seems that the prospect of acquiring knowledge seems to correlate with our intuitions about when it would be rational to believe.

#### 4 | CHALLENGING CLOCKS

If, as Bird suggests, knowledge is the standard against which we evaluate belief, the pattern of intuitions considered thus far makes sense. If epistemic norms concerned the truth and nothing but, our intuitions don't make sense. We think Gnosticism helps us identify the weak premise in the argument for Comparative Accuracy. The problem was Veritism, not Comparative Desirability. We find Comparative Desirability comparatively plausible. Think about the Lockean treatment of the preface. We assemble a large collection of similarly supported beliefs, one that has an exceptionally good ratio of true to false beliefs (the *tf-ratio*). The Lockeans, because they identify the *tf-ratio* with the ratio of desirable to undesirable belief (the *du-ratio*), can say that each belief in the collection is overwhelmingly likely to be one that's desirable from the epistemic point of view and *that* makes it difficult to deny that believing is rationally preferable to suspension. We can say the same thing if we identify the *du-ratio* with the ratio of knowledge to ignorance (the *ki-ratio*).<sup>17</sup> Each belief in the collection is also overwhelmingly likely to constitute knowledge. For us, this is what matters.

We see two patterns of inference as quite compelling. When it's nearly certain that some specific belief would be desirable from the epistemic point of view, it's hard to make the case that we should suspend. (To feel the force of this, it can help to abstract away from substantive specifications of what desirability amounts to.) When, however, it's nearly certain that some specific belief would be undesirable, it's hard to make the case that we shouldn't suspend. If we plug in our preferred conception of epistemic desirability, it's not surprising that we suspend in the lottery or the Moorean case. Even if the *tf-ratio* is good, the *ki-ratio* will be terrible. This is why there's an important difference between the lottery and the preface. Whilst we might be certain in the preface case that there's some flawed belief in the collection, what seems to matter to questions about whether to believe this or answer this question this way is how confident we can be that this particular attitude is desirable.

We should acknowledge one thing. We've said little in this paper about why our preferred understanding of epistemic desirability is better than the alternatives. Imaginative readers might find some alternative understanding of desirability and wonder why we should prefer our view to this one. One reason to prefer our approach is that it vindicates our intuitions. We should also note that our view of desirability is supported by three observations that strike us as quite plausible. The first is that a false belief is objectively epistemically undesirable. The second is that knowledge is more desirable from the epistemic point of view than mere true belief. The third is that if we're systematically disconnected from reality because, say, we're in Nozick's (1974) experience machine, we cannot have beliefs about the reality 'behind' the appearances that would be objectively epistemically desirable. The second point is compatible with the view that each true belief would be more desirable than suspension even if less desirable than knowledge. That view, when combined with Comparative Desirability, leads right back to the view that we should believe in the lottery case. That view, however, is ruled out by the third observation. There's nothing in Nozick's experience machine or the brain in the vat examples that ensures that none of our beliefs are *true*.<sup>18</sup> They can be true. In that scenario, however, where there's the occasional 'mere match' between belief and reality we don't think that the odd accurate belief has the properties that make desirable beliefs desirable. Gnosticism accommodates these three observations straightforwardly

given the conservative assumption that there's more to knowledge than mere true belief and it seems the most natural way to accommodate these three claims.

Isn't this good news for Knowledge Duplication? It's *goodish*. It's easy to see why someone might prefer Knowledge Duplication to Normality or the Lockean view, but reflection on ki-ratios points us towards an alternative knowledge-centred view. Consider four observations:

O1. It's not rational to believe outright: Alice attacked the guard, but we couldn't know whether she did.

O2. It's not rational to believe outright Alice attacked the guard if it's rational to believe outright that we couldn't know whether she did.

O3. If you and a peer were given different quiz questions and you're equally confident that you each knew the answer, you cannot rationally believe your answer if you learn that one of you didn't know the answer.

O4. If you and your peers were each given 200 questions and you learned that each of you made a mistake or two on your exams, you can each continue to rationally believe each of your answers having learned this.

These observations concern some Moorean cases, preface cases, and cases of peer disagreement. The cases differ in a myriad of ways, but what they share in common is some salient negative epistemic appraisal. In the Moorean cases, these appraisals target specific beliefs. In the preface and peer disagreement cases, they needn't. We need a theory that explains why rationality might be robust enough to persist in the face of information that supports some negative appraisals and not others. We don't think Knowledge Duplication is particularly helpful in thinking through these differences.

Knowledge Duplication makes quick work of O1. You cannot be sufficiently internally similar to someone who knows  $\langle p \ \& \ \sim Kp \rangle$  because knowledge is factive and distributes over the conjunction. Hence, Knowledge Duplication predicts that it's not rational to believe this conjunction.

How does Knowledge Duplication handle O2? This observation is about the conjuncts. Whilst it's obvious that no single thinker can know both conjuncts, it's not obviously impossible that Alice might be similar internally to Binh who knows  $p$  and internally similar to Charis who knows she doesn't know whether  $p$ .

Here are two bad suggestions for how we might account for O2 if we accept Knowledge Duplication. It would be bad to appeal to the conjunction rule (i.e., the idea that it's rational to believe conjunctions of things it is rational to believe) to explain O2.<sup>19</sup> It would also be a bad idea to insist that the set of things it's rational for Alice to believe must each be known by some single thinker at once. Remember our observation O4. To get the preface case right, we must reject the conjunction rule and we must allow that it's rational for Alice to believe inconsistent sets of propositions.

To explain O2, the defenders of Knowledge Duplication have to say that Alice cannot be sufficiently similar to two thinkers, one who knows  $p$  and another who knows she doesn't know whether  $p$ . They might be right to say this, but we see no justification for this assumption. We don't offer O2 as a counterexample. We see it as a data point and don't see how Knowledge Duplication predicts it.

Consider O3. Binh answers a question about bears. Charis answers a question about cheetahs. They know each other to be peers and learn that precisely one of them didn't know the answer to the quiz. Intuitively, they should suspend, not believe. Is it obvious that neither of them could be similar to someone who knew the relevant answers? That's not obvious to us. Ignoring the quizmaster's testimony, each could have been similar to knowers. When we take account of that testimony, is it obvious that *now* neither is similar to a knower? Again, it's not obvious to us. If we introduce Dinesh and his answer about Danish designers and reformulate the case so that Binh, Charis, and Dinesh are each .333 confident that they didn't know the answer, are they *now* sufficiently internal to some possible knower? How should we think about similarity relations and the ratio of happy to unhappy outcomes? We don't know. We just know that if the numbers get large enough, it's important to say that the thinkers can rationally believe. We feel like our intuitions aren't triggered by our understanding of the similarity relations between thinkers and possible knowers. We worry that that ignorance is best explained by the fact that there's no way to recover the relevance of the different ki-ratios from facts about these similarity relations.

We think the most natural explanation of O1-O4 gives pride of place to ki-ratios. Why can't you rationally believe in the original Moorean case? Because it's *certain* you cannot know the conjunction. That gives us O1. Why can you rationally believe in the preface cases? Because each belief is nearly certain to be something that can be known. That gives us O4. Why can't you rationally believe the conjuncts individually that make up the Moorean conjunction? Because these beliefs aren't sufficiently likely to be known. This predicts O2 and O3.

Think back to the quiz from the previous section:

Q1: What's the capital of California? Sacramento.

Q2: In the War of the Roses, which house battled the Yorks? The Lancasters.

...

Q200: Who was the first woman to win the Nobel Prize? Marie Curie.

We can imagine variations on this case where we play around with different information about the ki-ratios. If we assume that you're equally confident that you knew the answers initially, think that the likelihood of knowing the answer to some question is independent from that of knowing others, and we imagine increasingly bad epistemic news, we should reach some point where the quiz master says that you knew *some* answers but you nevertheless recognise that you are rationally required to suspend belief on Q65, the one about Irish neutrality in World War II. When you hit that point, the point at which you should suspend belief on whether Ireland was neutral, you should suspend on the other answers, too. Remember the similarity of support.

Here's an important detail. The quizmaster said that you knew some of the answers. Could you continue to know now if we agree that you've reached a point where rationality requires you to suspend because the ki-ratio is so poor and because in each case you'd be too likely to answer without knowing the answer? Remember that the questions concern a diverse range of topics. That you could have easily been wrong in answering that question about Ireland gives us little reason to think you could have easily been wrong in answering that question about California or the War of the Roses. Maybe it's both not rational to believe even if you assign positive probability to the hypothesis that you're in a position to know. If this is right, there might be a notion of risk

(i.e., the risk of believing what's not known) that's understood in terms of available information that bears on what's rational to believe and a notion of danger where objective features of the situation that make it dangerous to believe can preclude knowledge. That the risk of believing what's not known is high doesn't guarantee that there's a real danger in, say, answering Q134 incorrectly. It might not be rational to give or stick by the original answer *even if* it's possible in spite of the risks of not knowing that you're in a position to know.

We think rational thinkers will respond differently in the quiz cases depending upon how good or bad the ki-ratio is. We think that that's because these thinkers will conform to this norm:

**Minimisation:** It's rational to believe iff the expected objective undesirability of belief isn't greater than that of its alternatives.

Given Gnosticism, this gives us:

**Probable Knowledge:** It's rational to believe iff it's sufficiently likely given the thinker's information that she can know.<sup>20</sup>

We would add two claims to this:

**Not Worse:** Whilst it's undesirable to believe without knowing and undesirable to fail to believe what can be known, the latter isn't more undesirable than the former.

And:

**Not Much Worse:** Whilst it's undesirable to believe without knowing and undesirable to fail to believe what can be known, the former isn't vastly more undesirable than the latter.

This yields a view on which rationally believing requires that it's more likely than not that we can know and allows that we can rationally believe when it's not certain that we know. It handles all the cases that Knowledge Duplication can handle. In addition, it handles O1-O4 with ease.

## 4.1 | Clocks, Finally

We think that if Knowledge Duplication is going keep pace with Probable Knowledge, its proponents need to explain why there should be some connection between ki-ratios, the risk of believing what's not known, and the similarity relations that are supposed to matter to rational belief. It seems that they'd need to assume this to explain O1-O4:

**No Improbable Knowledge:** It's possible to know  $p$  only when it's more likely than not that the thinker can know  $p$ .

If it's possible to know when it's not likely that we know, it's not clear why O2 or O3 would hold. Because everyone who knows is maximally similar to themselves, Knowledge Duplication implies:

**Knowledge Sufficiency:** It's only possible to know  $p$  if it's rational to believe  $p$ .<sup>21</sup>

Given No Improbable Knowledge, it's only rational to believe when each of the conditions necessary on knowledge are likely to obtain. Thus, given Knowledge Sufficiency, the defenders of Knowledge Duplication would have to defend this, too:

**No Improbable Rationality:** It's only possible to rationally believe  $p$  if it's more likely than not that it's rational to believe  $p$ .

We think there are counterexamples to No Improbable Knowledge and No Improbable Rationality. The defenders of Knowledge Duplication are caught in a bind. If they don't embrace No Improbable Knowledge, they cannot explain O1-O4. If they do, their view is subject to counterexamples that won't arise for Probable Knowledge. Our view doesn't commit us to Knowledge Sufficiency or to No Improbable Knowledge.

Recall Williamson's (2014) clock. Williamson's clock has a face without markings and a minute hand that might be found at one of 60 positions. You know the clock is perfectly reliable. Let's suppose that you know you can tell the hand's location by looking within a margin of error ( $\pm 1$  minute) so that on the assumption that it's 2 minutes past the hour, the strongest thing you can know is that the time is within [1-3]. You also know that if it is 1 past, the strongest thing you can know is that the time falls within [12-2]. If it's 3, the strongest thing you can know is that the time falls within [2-4]. On Williamson's description of the case, when it's 2, you assign positive probability to three hypotheses compatible with your knowledge—that the strongest thing you can know is that the time falls within [12-2], that the strongest thing you can know is that the time falls within [1-3], or that the strongest thing you can know is that the time falls within [2-4].

Given this description of the case, we have a counterexample to No Improbable Knowledge. The strongest thing you can know (i.e., that the time falls within [1-3]) is known, but it's not more likely that you know this than you know something else.

What about No Improbable Rationality? Given your information, you assign equal probability to three hypotheses about what you can know and rationally believe given suppositions about the time. We list some of the propositions that might be rationally believed or known along the top of the columns and the times we might suppose it to be in the cells to the left of each row:

	t is in [12-2]	t is in [1-3]	t is in [2-4]	t is not 12	t is not 3	t is not 4
Time is 1	K, R	$\sim K, \sim R$	$\sim K, \sim R$	$\sim K, \sim R$	K, R	K, R
Time is 2	$\sim K, \sim R$	K, R	$\sim K, \sim R$	K, R	$\sim K, \sim R$	K, R
Time is 3	$\sim K, \sim R$	$\sim K, \sim R$	K, R	K, R	$\sim K, \sim R$	$\sim K, \sim R$

We'll suppose the time is 2 minutes past. Had you seen the clock at 1 minute past, you could rationally believe the time is within [12-2], but not that the time was within [1-3]. Had that been rational to believe as well, you could combine that belief with the belief that the time is not at 3 to work out that the time was either 1 or 2. Since you know you cannot know anything that strong, we know that you cannot rationally believe the time was within [1-3] if the time is 1. Notice that for each time it might be, there will be a proposition that's rational to believe (by virtue of Knowledge Sufficiency and the description of the case) where it's not likely given your information that it's rational to believe that thing. For example, it's not rational to believe the time lies within [12-2] if

the time is either 2 or 3, so even if the time is 1, that the time lies within [12-2] is not more likely than not to be rational given your information. Similarly, if the time were 1, it would be rational to believe the time wasn't 3, but it's not rational to believe that in most of the possibilities compatible with your information.

Williamson's preferred description of the clock case is controversial (Cohen & Comesaña, 2013; Goodman, 2013). We don't want to settle the controversy here, but we think it's fair to use it to make this point. If we think of Knowledge Duplication and Probable Knowledge as alternative developments of a knowledge-first approach to rational belief, it's fair to test these theories by thinking about what they'd say given possible discoveries about the nature of the knowledge relation. Williamson's take on his case give us some reason to think that we might discover something that seemed plausible in light of Radford's (1966) examples which is that a thinker can rationally doubt that they can know something and still be in a position to know. When, after much prodding and protests that she doesn't really know much about history, Alice tells us that the year William the Conquerer landed was 1066, it's very tempting to say that Alice knew after all. Given the information available to her, maybe she should be nearly certain that she did not know the answer. Probable Knowledge doesn't force our hand and it explains our conflicting responses to the case. In clear cases of knowledge, things look very good from the subject's point of view. When the information available indicates that it's not likely that something is known, we think the (correct) judgment that it's not rational to believe makes it harder for us to think of the case as a case of knowledge, but the judgment that it's not rational to believe doesn't make it *certain* that they don't know.

Our picture is something like this. Lots of things dispose us to believe (e.g., the senses, our apparent memories, hunches, stored beliefs, patterns of reasoning) and dispose us to doubt. Dispositions to suspend should be resisted when it's sufficiently likely given our information that we'll lose knowledge by suspending. Think of the preface. Dispositions to believe should be resisted when it's insufficiently likely given our information that by believing we'll know. Think of seeing \ures in the mist when searching for a friend in a crowded market. Could someone with the hazy impression knowingly identify a friend? Maybe. Should we take the risk because of this possibility? Maybe not.

We should revisit the rationale for Knowledge Duplication:

The idea behind [Knowledge Duplication] is that a subject's belief is justified just in case her intrinsic state is consistent with her having knowledge. It is trivial on ... [this view] that all cases of knowledge are cases of justified belief. So ... [this view] captures the right-to-left direction of  $J = K$ , and thus meets a weak version of the 'not-too-stingy' desideratum ... But it does not accept the left-to-right direction, which ran so starkly against the New Evil Demon intuition. Subjects in radical sceptical scenarios are among those who have justified beliefs that fall short of knowledge. My deceived counterparts with misleading experiences may be intrinsic duplicates of me; if so, their perceptual beliefs are justified, because mine is knowledge ... [This view] also places plausible constraints on justification: intuitively unjustified beliefs are typically such that the subject's internal state precludes knowledge (Ichikawa 2014: 190).

In this passage, Ichikawa explains why he thinks Knowledge Duplication builds on the strengths of  $J = K$  (i.e., the view that we justifiably believe iff we know) but improves upon it by capturing an important internalist intuition.<sup>22</sup> There's much here that we agree with. We agree, in

some sense, that if someone is internally similar to a rational believer, it shouldn't be that it's not rational for them to believe just because they're in a bad case. If rationality is a normative notion, it's properly better to think of it as connected to a kind of *subjective* normativity according to which the subject's information or perspective matters to normative standing and obscure objective matters of fact do not. If that's our interest, however, then whilst we agree with Ichikawa that the mere fact that someone doesn't know doesn't bear directly on whether it's rational to believe, we wonder why the mere fact that someone knows should bear directly on whether it's rational to believe. Objective undesirability doesn't make it subjectively or prospectively improper to believe. That's the lesson of the BIV case. By the same token, objective desirability doesn't seem to suffice for subjective or prospective propriety. That seems to be the lesson of the unmarked clock case. When we're internally like a knower *and* it's virtually certain given our evidence that we don't know, it seems unwise to take the risk of believing in ways that would be objectively undesirable. We think the intuitive rationale for Knowledge Duplication seems quite compelling when cases of improbable knowledge aren't salient. If a bit of autobiography can be forgiven, our intuitions flipped when the improbable knowledge cases became salient. Replacing modal considerations with probabilities and expectations seemed to us to capture the spirit of Knowledge Duplication but handle a wider range of intuitions.

We think we can further weaken the intuitive case for Knowledge Duplication by thinking about the relationship between knowledge and criticism. Inquiry, we think, is a practical endeavour, but one worth thinking about because of the aim to acquire an epistemic good (e.g., coming to know the answer to a question). Suppose Alice has access to two sets of encyclopaedias, the 14th and 15th edition. The main difference between them is that a handful of errors found in the 14th were corrected so the ki-ratio of the 15th edition is slightly better. Seated between the two sets, we'd be hard pressed to think of any reason why Alice should prefer using the 14th instead of the 15th to answer her questions about marsupials, but that's what she did. She thus took an unnecessary risk of relying on one of the false entries that was corrected in the 15th edition. She could still acquire the knowledge she sought (e.g., to know whether all marsupials are native to Australia), but we can still say that even if she knows because of what she read that there are marsupials endemic to North America that she took an unnecessary risk. We don't think that the happy outcome of having acquired knowledge speaks to the question whether it was rational for her to conduct her inquiry like this. This isn't a pure epistemic case, but we think that it suggests that at least sometimes being internally similar to a knower isn't a defence against the charge that our responses are irrational because we've taken the wrong risks in the pursuit of something epistemically desirable. We're open to the possibility that something similar can happen in the epistemic case, that there might be knowledge available to us that we won't take up if we're reasonable (Lasonen-Aarnio, 2010).

## 5 | JUSTIFICATION BY DEGREES

Brown (2018) objects to the proposal that justified beliefs are just items of knowledge on the grounds that it tells us nothing interesting about *graded* justification.<sup>23</sup> If we identify knowledge with justified belief, what would it mean to say that some belief on this view is "more justified" than another? If I mistakenly believe *p* and you mistakenly believe *q*, neither of us know. We might still want to say that you were more justified in your beliefs. Saying that a belief is justified iff known sheds no light on how this might be. Similar worries might arise for Knowledge Dupli-



cation. We don't know how to give a plausible account of the degree of justification in terms of similarity relations between thinkers and possible knowers.

This line of objection to  $J = K$  might be too quick. Couldn't Williamson simply say that the degree of justification corresponds to something like evidential probability? Wouldn't that give Brown what she seeks? The problem with this response is that we want people to say *plausible* things about degrees of justification. Assuming that degrees of justification are theoretically important *and* that there are important connections between degrees of justification and justification simpliciter, we want a theory that predicts that Binh cannot have more justification than Alice if only Alice's beliefs are justified. As we've seen in our discussion of Comparative Accuracy, we know that if degrees of justification are understood in terms of evidential probability, we'll get violations of this principle:

**Comparative Degrees:** If it's rational for you to believe  $q$  and the degree of justification you have to believe  $p$  is greater, it's rational to believe  $p$ .

Our impression is that part of the reason Brown is troubled by Williamson's view is that it seems it won't vindicate Comparative Degrees.

If a theory of justification should accommodate Comparative Degrees, our three main theories don't fare terribly well. The Lockeans can give us a nice theory of the relation between justification simpliciter and degree of justification. On their view, the degree of justification can be understood in terms of the probability of the truth of the target proposition. It predicts that beliefs are justified simpliciter when the degree of justification crosses a threshold. Alas, their view of justification simpliciter conflicts with our intuitions. We don't see crossing that threshold as sufficient for justification. This is one lesson of the lottery. The defenders of Normality can offer an alternative understanding of degree of justification by offering a theory of degrees of normic support (Smith, 2016). We can have justification simpliciter in preface cases. Normality predicts we lack justification simpliciter in such cases, so this doesn't give us what we want, either. If Normality says we don't have a sufficient degree of justification, that's not satisfying. If we have beliefs that are justified simpliciter, we couldn't lack the necessary degree of justification for having beliefs that are justified simpliciter. If, on the other hand, Normality says that we have a high degree of justification in the preface cases but lack justification simpliciter, their view will inevitably violate Comparative Degrees. We can cook up a non-preface case in which the degree of justification is lower but we have beliefs that are justified simpliciter.

We're off to a bad start. Let's make things worse. Degree of justification should fit with our best understanding of what makes objectively epistemically desirable beliefs desirable. Here's the rough idea. Rational thinkers shouldn't be indifferent to the presence or absence of sufficient rational support. If, however, there were no interesting connection between degree of justification and epistemic desirability, it would seem that they could be rationally indifferent to its presence. Suppose that believing  $p$  had a greater degree of justification than believing  $q$  and that believing  $p$  was less likely to be desirable than believing  $q$ . If believing  $q$  were rational, Comparative Degrees tells us that it should be rational to believe  $p$ , but for all that's been said, suspending on whether  $p$  might still have a greater degree of expected desirability than believing. We want to rule that sort of thing out.

The easy fix is to connect degree of justification to the expected desirability of belief. We can accommodate Comparative Degrees and Comparative Desirability if we identify the degree of justification with the probability of meeting the conditions necessary for knowing. This should solve the problem by ruling out the possibility that rational thinkers will be rightly indifferent to

the presence or absence of rational support, yielding a theory of rationality that's extensionally adequate and that seems to give us the right verdicts in the right ways.

## 6 | A COMPELLING IDEA?

We've proceeded as if theories of rational belief are supposed to tell us what's epistemically desirable and then identify principles linking objective desirability to rationality. Our preferred theory of desirability is Gnosticism. Our preferred theory of the connection between rationality and desirability is Minimisation. Given its popularity in the practical domain, it seemed only reasonable to give it a try in the epistemic domain.

We think it's a helpful exercise to see what happens if we try to formulate rival theories in similar ways (i.e., by trying to see how some theory might understand objective desirability and its connection to rationality). It would be worrisome, we think, if we cannot reconcile some view with a plausible theory of desirability and a plausible theory of the connection between desirability and rationality.

We don't know where the defenders of Normality stand on Veritism, but their view is equivalent to the conjunction of Veritism and this principle linking rationality and desirability:

**Normally Desirable:** It is rational for you to believe if believing is objectively epistemically desirable in the most normal worlds compatible with your evidence.

The defenders of Normality don't have to say that this is a genuine principle of epistemic rationality. We worry that theories that don't posit principles linking rationality to desirability have less explanatory power than those theories that do. Is it supposed to be a brute fact that normic support is the stuff epistemic rationality is made of?

One reason to be mildly sceptical about Normally Desirable is that it conflicts with this:

**Better Chances:** When considering two lotteries, L1 and L2, L1 is more desirable than L2 if they have the same two possible outcomes but it's more likely that the more desirable outcome would result from choosing L1.<sup>24</sup>

To appreciate the conflict, think about the risks in a testimony case (L1) a preface case (L2). If we're assuming Veritism, we can arrange things so that we're more likely to get the desirable outcome in the preface case. Normally Desirable says that regardless of whether we understand desirability in terms of accuracy or knowledge, believing is preferred to suspending in the testimony case but suspending is preferred to believing in the preface case. If we think of 'choosing' L1 or L2 as believing and suspending as not choosing, Normally Desirable says that we should choose L1 and not choose L2. If the defenders of Normality want to accommodate Better Chances, they could reject Veritism and claim that normic support is the stuff objective epistemic desirability is made of, but we think this isn't a plausible claim about objective epistemic desirability. Everyone seems to agree that false beliefs are undesirable from the epistemic point of view.

We find weak inductive evidence against Normality when we examine the practical analogue of Normally Desirable:

**Normally Desirable+:** It is rational for you to choose an option if there isn't an option that's more desirable in the most normal worlds compatible with your evidence.

Suppose that if you take L1, you'll win \$5 if  $p$  and nothing otherwise. If you take L2, you'll win \$5 if  $q$  and \$.50 otherwise. Suppose further that you're equally confident in  $p$  and in  $q$  and have normic support for believing  $p$  and for believing  $q$ . It seems irrational to choose L1 over L2, but Normally Desirable+ doesn't predict this. In the most normal worlds compatible with your evidence, the outcomes of both bets would be \$5, but it's hard to believe that the sweetener attached to the second bet doesn't make L2 the option to choose.

Consider Knowledge Duplication. Knowledge Duplication is equivalent to the conjunction of Gnosticism and this:

**Optimality Duplication:** It's rational to believe when you're sufficiently similar on the inside to someone who believes and therein maximises objective epistemic desirability.

This principle, considered abstractly, seems strange. If we're right that a thinker can know some proposition in a collection of similarly supported propositions despite her knowledge that the *ki*-ratio is bad, why would we think that it's rational to continue believing the propositions that happen to be known when there's so little to be said for the similarly supported propositions that are neither known nor likely to be known? It's also strange because the principle seems completely misguided if we combine it with alternative conceptions of epistemic desirability. Combine Optimality Duplication with Veritism and we're left with the view that it's rational to believe each ticket in a large lottery will win. Even though it's virtually certain that you'll believe falsely by believing such things, there's someone in modal space who takes the chance and forms a true belief. We doubt this principle identifies a normatively significant relation between options and the properties that make them desirable. Defenders of Knowledge Duplication aren't committed to Optimality Duplication, but it would be worrisome if we couldn't see Knowledge Duplication as combining a good theory of desirability with a good theory of the connection between desirability and rationality.

We find weak inductive evidence against Optimality Duplication by thinking about the practical analogue the principle:

**Optimality Duplication+:** It's rational to choose X out of some option set when you're sufficiently similar internally to someone who chose X out of this option set and thereby maximised objective desirability.

If you're sufficiently similar on the inside to some adventurous internet shopper who paid millions for a lottery ticket advertised as 'the winning' ticket on eBay and won many millions more, Optimality Duplication+ says that it's rational for you to follow their lead. This is deeply implausible. We cannot think of any reason why we should be optimistic about the epistemic analogue of this principle.<sup>25</sup>

Some defenders of Knowledge Duplication might be sympathetic to Veritism.<sup>26</sup> Knowledge Duplication follows from the conjunction of Veritism and this principle:

**Known Optimality Duplication:** It's rational to believe when you're sufficiently similar on the inside to someone who *knows* that they will maximise objective epistemic desirability by believing.

We move knowledge out of the characterisation of the desirable epistemic result and find space for it in a principle linking desirability to rationality.

Consider the practical analogue of this principle:

**Known Optimality Duplication+**: It's rational to choose X out of some option set when you're sufficiently similar on the inside to someone who *knows* that they will maximise objective desirability by choosing X out of this option set.

This principle is an improvement upon Optimality Duplication+ (e.g., it predicts that it's irrational to pay millions for lottery tickets that possibly will yield millions more), but Known Optimality Duplication+ isn't plausible. Consider Parfit's (1988) miners case. We know that all the miners are in one of two shafts, but have no more reason to think they are in shaft A than shaft B. It starts to rain. If we leave the shafts uncovered, water will partially fill the mines killing one miner in whichever mine the team is in. If we cover one shaft, we'd save each miner in that shaft but completely flood the other and that would kill every miner in the shaft. It's rational to leave the shafts uncovered, but we *know* that that's the only response that doesn't maximise objective desirability. Counterexamples to Known Optimality Duplication+ aren't counterexamples to Known Optimality Duplication, but they give us some reason to pause. If we love miners, shouldn't we be guided by probabilities rather than knowledge? If we loved the truth, doesn't something similar hold?

It's rare for theories of rational belief to be stated explicitly in terms of a theory of objective desirability and a principle that explains the connection between objective desirability to rationality. Boghossian is an exception. He's struck by the analogy between the imperatives to believe the truth and shun error and the trader's attempts to buy low and sell high (2008: 101). This analogy isn't perfect (e.g., you'd hope someone looking after your investments wouldn't be put off by trade-offs), but it's not terrible. Still, some readers might be put off by the apparently consequentialist character of our way of framing things.

What can we say in defence of our approach? First, judge the approach by its fruits. We see no alternative that predicts the right verdicts. Second, consequentialism seems compelling. Foot attributes this to the thought that it cannot be right to prefer a worse state of affairs to one that's better (1985: 198). If we think reasonable people won't prefer what's worse, by their lights, to what's better, reasonable people confident in their normative outlook will prefer theories that predict that the thing to do or the thing to believe will pattern with our understanding of what's desirable. In a setting where nobody defends the view that the thing to do or believe is determined by what's *objectively* most desirable, it's natural to look for views that use some information-relative notion of desirability to vindicate our intuitions about choice or about belief. That's what we've done. Our view seems compelling because it captures the idea that made consequentialism seem compelling.

Foot didn't deny this idea that reasonable people don't prefer what's worse by their own lights, but Foot did deny consequentialism. She thought that if we thought of the desirable in the way reasonable people do, we'd see that the compelling idea didn't support consequentialism over non-consequentialism. We feel that our approach is compatible with non-consequentialism even if it can be presented in a (nearly) consequentialist way. Dreier (2011) reminds us that we can make all kinds of views look (nearly) consequentialist. We can speak *as if* we think it's a bad state of affairs that people believe without knowing if we think that desirability should be understood in terms of some kind of goodness, in terms of fittingness, or in terms of norm conformity. We see the flexibility of the approach as part of its appeal.

If we thought that knowledge was a good thing to exist and ignorance was a bad thing, we *could* interpret our view as saying that Probable Knowledge holds because conforming to Probable Knowledge ensures that we believe in ways that maximise expected epistemic value. Speaking in our way, however, doesn't commit us to anything as controversial as the view that epistemic reasons are value-based.<sup>27</sup> We could say instead that there are epistemic norms that enjoin us to believe things we can know (K+) and to refrain from believing what we cannot know (K-) and that what makes it objectively desirable or undesirable to believe or suspend has to do with conforming to such norms. As with the value-based theory, we don't want to say that there's a direct connection between rationality and objective desirability in this sense. Instead, we'd want to say that when we assign positive probability to the hypothesis that these objective reasons obtain, their rational strength is determined by their objective weight and the probability that these reasons are actual. We would break with recent authors who suggest that it's by being in a position to know facts that they have a bearing on rationality (Kiesewetter, 2017; Lord, 2018) and suggest instead that these norms always bear on the rationality of belief and suspension when we assign positive probability to the possibility that we'd violate them. It's a familiar proposal now in the literature on how non-consequentialists should think about uncertainty that given the (objective) duty to conform to non-epistemic norms that the force or weight of the objective reasons to conform to these norms should be discounted probabilistically (Lazar, 2020; Olsen, 2018). If we believed in epistemic norms of the kind just described above, we could say that Probable Knowledge is true because we minimise expected objective wrongfulness if we believe in accordance with it.

We don't propose to 'consequentialise' epistemology, but we think we can come close by teleologising it. Believers can be described *as if* they want to inscribe something of interest into the belief box on the condition that it's known. When it can be known, the failure to inscribe it is objectively undesirable to some degree. When it cannot be known, the failure that results from inscribing it is objectively undesirable to a greater degree. When it comes to questions of whether it's rational to believe or not, we compare the alternatives of believing and not believing in terms of the expected undesirability of each.

Dreier conjectures that every *plausible* ethical theory can be consequentialised. This raises an interesting question about the status of the debate between consequentialists and non-consequentialists in ethics. In epistemology, it might be that *no* plausible theory can be consequentialised. Still, we think something close to this idea should be considered. In consequentialist ethical theories, we list the feasible alternatives and rank them in terms of some notion of desirability that takes account of the total state of the world if some alternative is chosen. In evaluating attitudes, we seem to care about the properties of this attitude in isolation. Isolating this attitude from others doesn't prevent us from stating the conditions under which it is desirable or undesirable to hold it or use comparisons of desirability to give us a notion of expected undesirability.

## 7 | CONCLUSION

We find the right theory of rational belief by combining the right theory of objective epistemic desirability with a conservative theory of the connection between objective desirability and rationality of response. If we combine a theory of desirability that differs only slightly from Veritism with the idea that we should minimise expected objective desirability, we arrive at the only view that gets our cases right. It predicts that the preface differs from the lottery. It predicts that the ratio of success to failure might bear on rational status without bearing on knowledge, thus explaining

why rational belief needs more than the mere possibility that we might know against all expectation. It predicts that rational beliefs are both more likely to be desirable than irrational beliefs and that they're justified to a greater degree than irrational beliefs. These, we think, are two truisms that surprisingly aren't accounted for by most familiar theories. We propose that rational beliefs are rational because they are promising moves to make in the pursuit of knowledge.

## ACKNOWLEDGEMENTS

Open access publishing facilitated by Australian Catholic University, as part of the Wiley - Australian Catholic University agreement via the Council of Australian University Librarians.

## ORCID

Julien Dutant  <https://orcid.org/0000-0002-2939-5670>

## NOTES

<sup>1</sup>For their helpful feedback we would like to thank Maria Alvarez, Robert Audi, Rima Basu, Bob Beddor, Alexander Bird, Kyle Blumberg, Darren Bradley, Bill Brewer, Jessica Brown, J. Adam Carter, Sam Carter, Tez Clark, Nilanjan Das, Kevin Dorst, Stephen Finlay, Branden Fitelson, Carolina Flores, Rachel Fraser, Dmitri Gallow, Sandy Goldberg, Simon Goldstein, Jeremy Goodman, John Hawthorne, Nick Hughes, Jessica Isserow, Liz Jackson, Chris Kelp, Benjamin Kiesewetter, Jennifer Lackey, Maria Lasonen-Aarnio, Kathryn Lindeman, Artūrs Logins, Errol Lord, Winnie Ma, Aidan McGlynn, Eliot Michaelson, Jennifer Nagel, David Papineau, Francesco Praolini, Richard Pettigrew, Sven Rosenkranz, Bernhard Salow, Susan Schellenberg, Mark Schroeder, Matty Silverstein, Mona Simion, Martin Smith, Ernest Sosa, Margot Strohminger, Kurt Sylvan, Sergio Tenenbaum, Jonathan Way, Ralph Wedgwood, Robbie Williams, Christopher Willard-Kyle, Tim Williamson, Raphael Woolf, Juhani Yli-Vakkuri, as well as audiences at Barcelona, Glasgow, the Dianoia Institute of Philosophy, King's College London, LSE, NYU-Abu Dhabi, Konstanz, and St. Andrews.

<sup>2</sup>Formulating the Lockean view is tricky because there are some cases where propositions are highly probable on the thinker's evidence provided that they aren't believed (e.g., the proposition expressed by claims like, 'It's raining but I don't believe that it is'). Thus, it's not clear whether it's better to formulate the view as saying that it's rational to believe when it's rational to be very confident that something is true or to say that it's rational to believe when it's rational to be very confident that some belief would be correct or accurate.

<sup>3</sup>For reasons of space, we abstract away from the differences between these approaches. We don't think the differences will matter here.

<sup>4</sup>We have been asked why we do not discuss (Conee & Feldman, 2004) evidentialist view, for example. In a way, we do. They think that a thinker's evidence determines what that thinker has justification to believe and seem to be open to proposals about how we might distinguish between the relations between evidence and belief that confer justification upon the latter and those that do not. We can see our three theories as telling us why some evidence does (and some evidence does not) justify belief. There are places where they seem to suggest that they favour something like (McCain, 2014) explanationist view of evidential support. We do not discuss that view here, but readers can decide whether they think there is an explanationist view that can incorporate our principles and get the data right. We explain why we think the explanationist view is *not* explanatorily adequate and why our view seems to capture what made that view seem attractive in (Littlejohn & Dutant, forthcoming).

<sup>5</sup>While Papineau (2021) argues that we shouldn't be unwilling to punish on the basis of naked statistical evidence, he agrees with the points about blame. That's sufficient for our purposes given that only outright beliefs about blameworthy actions rationalise attitudinal blame. We think it's helpful to follow the lead of Adler (2002), Gordon (1990), Owens (2013), and Unger (1975) in using the connections between rational belief and reasonable emotional response to try to identify the theory we're after.

<sup>6</sup>Notice there are two claims here. The first is that we know that we don't know in lottery cases. The second is that when we know we don't know in these sorts of cases, suspension is mandatory. For dissent on the first point, see Engel (2021), Gibbons (2013), and Sosa (2015). For dissent on the second, see McGlynn (2013).

<sup>7</sup>We explore the problems that Moorean cases present for Normality in (Littlejohn & Dutant, 2020).

- <sup>8</sup>Note that on views that say that we can know in lottery cases, it seems we should also be *highly* confident that we know. In turn, it seems that it should be reasonable to blame in Prisoners since it would be virtually certain that we knew the defendant to be guilty. That runs contrary to intuition.
- <sup>9</sup>See also Neta (2022) for a discussion of the importance of controlling the possible factors that can vary the threshold in evaluating the Lockean view. We assume that the threshold doesn't shift because of things like the number of beliefs we form and the increased risk of acquiring false beliefs. If the threshold shifts, it will have to do with the desirability of particular attitudes.
- <sup>10</sup>We are also assuming that, from the epistemic point of view, if it is undesirable to believe  $p$  (e.g., because  $p$  is false) and undesirable to believe  $q$  (e.g., because  $q$  is false), it is not more undesirable to hold one of these beliefs than the other. Presumably, if we thought that the contents of these beliefs and, say, their practical importance mattered to epistemic undesirability, we would not need the lottery cases to create trouble for Comparative Accuracy. We can ultimately relax this assumption and allow for variations in undesirability (e.g., that it might be worse to be wrong about  $p$  than it is to be wrong about  $q$  because, say,  $p$ -attitudes are important and  $q$ -attitudes are not), but then we would reject Comparative Desirability. We could still retain Minimisation (introduced later in the discussion) and our view would still emerge as the one view that gets the data right.
- <sup>11</sup>Here are some important examples from the literature. Neta (2022) approaches the lottery cases using defeaters and a broadly evidentialist framework that's neutral on Veritism. Leitgeb (2014) never rejects Veritism. He imposes a stability requirement to explain why rationally assigning a high degree of confidence alone isn't sufficient for rational belief. Nelkin (2000) approaches the lottery using a broadly explanationist view according to which it's not rational to believe things that are highly probable when we can see there's no explanatory or causal link between the target facts and our beliefs. Gardiner (2019, 2020) approaches these cases using a relevant alternatives framework on which rationality can require eliminating certain alternative possibilities when highly improbable without imposing a general threshold on rational belief across the board. See also Jackson's (forthcoming) discussion of the rational connections between belief and credence.
- <sup>12</sup>It can be difficult to distinguish gnostics (in this sense) from their critics because whilst many epistemologists will say that knowledge is the highest good or the aim of belief, some might accept the view, defended by Sosa (2015) that accurate beliefs that fail to constitute knowledge are less desirable than knowledge but more desirable than suspension. That view predicts, along with Veritism, that believing is preferable to suspending in lottery cases.
- <sup>13</sup>This is our twist on Makinson's (1965) preface case. The expert testimony is introduced to help remove the temptation to think that the thinker isn't convinced that an error has been made. The testimony should convince a reasonable thinker, we think. We also want to remove the complication that thinkers might not be able to survey their beliefs. For all that's been said, our thinkers have committed the questions and their answers to memory. Note that our quizmasters will often declare that people don't *know*. The quizmaster doesn't necessarily say that errors were made.
- <sup>14</sup>Note that a thinker might discover, having initially thought her beliefs were based indirectly on a testifier's observations, that her testifier had relied on naked statistical evidence. In this case, we think the thinker should suspend even though her confidence that the target proposition is true might increase. In our case, we don't think we find this sort of thing happening. In our case, there's an increase of confidence and the beliefs remain rational. We think neither the Lockean view nor Normality properly explains the surprisingly complicated relationship between reasons to suspend and reasons to increase confidence. Credit, though, to the defenders of Normality for pointing to this surprising phenomenon.
- <sup>15</sup>Granted, the separateness of propositions intuition has been contested. See Talbot (2014) and (Dunn & Ahlstrom-Vij, 2017).
- <sup>16</sup>Jackson (forthcoming) floats the possibility that statistical information that warrants a high degree of credence that our beliefs contain undesirable properties impacts our beliefs differently than information that warrants outright belief that our beliefs contain undesirable properties. This seems to fit with something we often see in discussions of the preface. In that literature, some have tried to suggest that preface writers aren't inconsistent because they don't believe that their work contains errors (Clarke, 2017). That might be true in the usual set up. It's one reason why we've introduced evidence of an error that comes from expert testimony. We think fixating on whether we believe outright or strongly suspect that a belief is false is unhelpful. Both can be toxic. We don't see any difference in kind between these toxins.
- <sup>17</sup>We are using 'ignorance' so that a thinker is 'ignorant' if he believes but doesn't know. Talk of 'botched knowing' might be clearer, but it's cumbersome.

- <sup>18</sup>See Kraft (2013) for an argument that most of us would know we held some true beliefs even if the sceptical hypotheses were true.
- <sup>19</sup>For discussions of the conjunction rule in the preface or the lottery, see Praolini (2019), Ryan (1996), Makinson (1965), and Smith (2022).
- <sup>20</sup>We think that one virtue of this view is that it doesn't need a no-defeater clause, provided that we understand defeaters as indicators that one or more of the conditions necessary for knowledge hasn't been met. In (Littlejohn & Dutant, 2021), we argue that defeaters are really indicators of ignorance. We get partial defeat when some information lowers the probability that we can know and full defeat when that probability is brought below threshold. This view explains the puzzling phenomenon of evidence that can both boost a proposition's probability whilst being rationally toxic.
- <sup>21</sup>Whilst this claim is widely accepted, the reasons for accepting it now differ quite significantly and pull in different directions. For some, it's a consequence of their preferred view of knowledge and seems to conflict with the idea that knowledge is the normative standard against which belief is evaluated (Schroeder, 2021). For some, it's supposed to be a consequence of the idea that knowledge is the normative standard against which belief is evaluated Sutton (2005), Bird (2007), Ichikawa (2014), Rosenkranz (2021), Gibbons (2013), Reynolds (2013), and Smithies (2012). We're sceptical. We think that non-human animals can know (e.g., our dogs and cats know that it's time for dinner, that we're home, etc.) but we doubt that there are attitudes that they should or shouldn't have. If they cannot be held responsible for their responses, it's not clear that the things they know are *rational* for them to believe, not if rationality is a normative status. See Kornblith (2002) and Sylvan (2018) for discussion. If knowledge doesn't essentially 'contain' any normative notion, it's eligible to be the non-normative notion that plays the kind of right-making role that accuracy plays in the Lockean theory.
- <sup>22</sup>For defences of J=K, see Sutton (2005), Williamson (2007), McDowell (1995), Littlejohn (2013). We don't feel the need to *reject* J=K in this paper because we think that if anyone takes J=K seriously, they're interested in something like *objective* justification. It's clear that theories of rationality don't aspire to be theories of that. Like Smithies (2012), we see this project as that of giving a 'subjective' notion of justification.
- <sup>23</sup>We don't discuss here some of the concerns that have been raised about the notion of degrees of justification. See Hawthorne and Logins (2021) for discussion.
- <sup>24</sup>This is one of the rationality conditions that Resnik (1987) uses to justify expected utility theory.
- <sup>25</sup>Smithies (2015) proposes something quite different. He thinks that we're justified in X-ing when we're justified in believing that X-ing would be correct. Given his assumption that beliefs are correct iff they constitute knowledge, this vindicates Knowledge Sufficiency and much besides. Here are two concerns with his crucial linking principle. What should we say about suspension? It's not clear that suspension *has* a correctness condition, but then the linking principle commits us to the view that suspension is never justified. We deny that. Suspension is justified when considering whether the number of stars is even. Suppose, however, that suspension has a correctness condition because, say, every incorrect response has a correct alternative or because correctness corresponds to objective propriety and it's objectively proper to suspend in this case. His linking principle still seems problematic. What if we're in a situation where we cannot justifiably believe we know and cannot justifiably believe we don't know? If there's no response we have justification to believe is correct, it seems we'd be stuck in a kind of dilemma. Smithies might respond to this by suggesting that this just shows that we should embrace a kind of access internalism that ensures that we always have justification to either believe we can know or cannot know. Obviously, that's a controversial view to defend. We note that the practical cases give us reason to doubt the practical analogue of the linking principle. In the miners case, we should choose an option we are *certain* won't be objectively best. That looks like a case where we know an action is justified *and* know that that that action is *not* correct. Our conjecture is that his view is one of many very similar views in the literature that are *very* close to getting everything right, but needs some readjustment to deal with uncertainty about whether we can know. We see our view as the next step in the development of this knowledge-first approach.
- <sup>26</sup>We don't know if Schroeder (2021) accepts Veritism. He defends Knowledge Duplication and seems to reject Gnosticism.
- <sup>27</sup>For scepticism about value-based approaches to epistemic reasons or rationality, see Côté-Bouchard (2017) and Raz (2011). See Lemos (1994) for a considered defence of the view that states of affairs in which we have true beliefs or knowledge can be good.



## REFERENCES

- Adler, J. E. (2002). *Belief's Own Ethics*. MIT University Press.
- Basu, R. (2019). What We Epistemically Owe to Each Other. *Philosophical Studies*, 176(4), 915–931. <https://doi.org/10.1007/s11098-018-1219-z>
- Basu, R., & Schroeder, M. (2019). Doxastic Wronging. In B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment in Epistemology* (pp. 181–205). Routledge.
- Berker, S. (2013). Epistemic Teleology and the Separateness of Propositions. *The Philosophical Review*, 122(3), 337–393. <https://doi.org/10.1215/00318108-2087645>
- Bird, A. (2007). Justified Judging. *Philosophy and Phenomenological Research*, 74(1), 81–110. <https://doi.org/10.1111/j.1933-1592.2007.00004.x>
- Boghossian, P. (2008). *Content and Justification: Philosophical Papers*. Oxford University Press.
- Buchak, L. (2014). Belief, Credence, and Norms. *Philosophical Studies*, 169(2), 1–27. <https://doi.org/10.1007/s11098-013-0182-y>
- Clarke, R. (2017). Preface Writers Are Consistent. *Pacific Philosophical Quarterly*, 98(3), 362–381. <https://doi.org/10.1111/papq.12130>
- Cohen, S., & Comesaña, J. (2013). Williamson on Gettier Cases and Epistemic Logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1), 15–29. <https://doi.org/10.1080/0020174X.2013.775012>
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Oxford University Press.
- Côté-Bouchard, C. (2017). Is Epistemic Normativity Value-Based? *Dialogue*, 56(3), 407–430. <https://doi.org/10.1017/s0012217317000506>
- Dorst, K. (2019). Lockeans Maximize Expected Accuracy. *Mind*, 128(509), 175–211. <https://doi.org/10.1093/mind/fzx028>
- Dreier, J. (2011). In *Defense of Consequentializing*. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics* (Vol. 1).
- Dunn, J., & Ahlstrom-Vij, K. (2017). Is Reliabilism a Form of Consequentialism?.
- Easwaran, K. (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>
- Engel, M. (2021). Lotteries, knowledge, and inconsistent belief: Why you know your ticket will lose. *Synthese*, 198(8), 7891–7921. <https://doi.org/10.1007/s11229-020-02555-w>
- Enoch, D., & Spectre, L. (2021). Statistical Resentment, Or: What's Wrong with Acting, Blaming, and Believing on the Basis of Statistics Alone. *Synthese*, 199(3–4), 5687–5718. <https://doi.org/10.1007/s11229-021-03042-6>
- Foley, R. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of Belief* (pp. 37–47). Springer.
- Foot, P. (1985). Utilitarianism and the Virtues. *Mind*, 94(374), 196–209.
- Gao, J. (2019). Credal pragmatism. *Philosophical Studies*, 176(6), 1595–1617. <https://doi.org/10.1007/s11098-018-1081-z>
- Gardiner, G. (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>
- Gardiner, G. (2020). Relevance and Risk: How the Relevant Alternatives Framework Models the Epistemology of Risk. *Synthese*, 199(1–2), 481–511. <https://doi.org/10.1007/s11229-020-02668-2>
- Gibbons, J. (2013). *The Norm of Belief*. Oxford University Press.
- Goodman, J. (2013). Inexact Knowledge without Improbable Knowing. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1), 30–53.
- Goodman, J., & Salow, B. (forthcoming). Epistemology Normalized. *Philosophical Review*.
- Gordon, R. M. (1990). *The Structure of Emotions: Investigations in Cognitive Philosophy*. Cambridge University Press.
- Harman, G. (1968). Knowledge, Inference, and Explanation. *American Philosophical Quarterly*, 5(3), 164–173.
- Hawthorne, J., & Logins, A. (2021). Graded Epistemic Justification. *Philosophical Studies*, 178(6), 1845–1858. <https://doi.org/10.1007/s11098-020-01512-0>
- Hyman, J. (1999). How Knowledge Works. *Philosophical Quarterly*, 49(197), 433–451.
- Ichikawa, J. J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, 44(2), 184–206.
- Jackson, E. (forthcoming). On the Independence of Belief and Credence. *Philosophical Issues*.
- Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford University Press.

- Kornblith, H. (2002). *Knowledge and its Place in Nature* (Vol. 115, Issue 2, pp. 246–251). Oxford University Press.
- Kraft, T. (2013). Sceptical Scenarios Are Not Error-Possibilities. *Erkenntnis*, 78(1), 59–72.
- Lasonen-Aarnio, M. (2010). Unreasonable Knowledge. *Philosophical Perspectives*, 24(1), 1–21. <https://doi.org/10.1111/j.1520-8583.2010.00183.x>
- Lazar, S. (2020). Duty and Doubt. *Journal of Practical Ethics*, 8(1), 28–55.
- Leitgeb, H. (2014). The Stability Theory of Belief. *Philosophical Review*, 123(2), 131–171.
- Lemos, N. (1994). *Intrinsic Value: Concept and Warrant*. Cambridge University Press.
- Littlejohn, C. (2013). The Russellian Retreat. *Proceedings of the Aristotelian Society*, 113, 293–320.
- Littlejohn, C. (2020). Truth, knowledge, and the standard of proof in criminal law. *Synthese*, 197(12), 5253–5286. <https://doi.org/10.1007/s11229-017-1608-4>
- Littlejohn, C., & Dutant, J. (2020). Justification, Knowledge, and Normality. *Philosophical Studies*, 177(6), 1593–1609.
- Littlejohn, C., & Dutant, J. (2021). Defeaters as Indicators of Ignorance. In M. Simion & J. Brown (Eds.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Littlejohn, C., & Dutant, J. (forthcoming). Knowledge and Prizes. In A. Logins & J.-H. Vollet (Eds.), *Putting Knowledge to Work: New Directions for Knowledge-First Epistemology*. Oxford University Press.
- Locke, D. T. (2014). The Decision-Theoretic Lockean Thesis. *Inquiry*, 57(1), 28–54. <https://doi.org/10.1080/0020174X.2013.858421>
- Lord, E. (2018). *The Importance of Being Rational*. Oxford University Press.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25(6), 205.
- McCain, K. (2014). *Evidentialism and Epistemic Justification*. Routledge.
- McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, 55(4), 877–893.
- McGlynn, A. (2013). Believing Things Unknown. *Noûs*, 47(2), 385–407. <https://doi.org/10.1111/j.1468-0068.2011.00843.x>
- Miracchi, L. (2015). Competence to know. *Philosophical Studies*, 172(1), 29–56.
- Miracchi, L. (2019). When Evidence Isn't Enough: Suspension, Evidentialism, and Knowledge-first Virtue Epistemology. *Episteme*, 16(4), 413–437. <https://doi.org/10.1017/epe.2019.34>
- Moss, S. (2018a). Moral Encroachment. *Proceedings of the Aristotelian Society*, 118(2), 177–205.
- Moss, S. (2018b). *Probabilistic Knowledge*. Oxford University Press.
- Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, 109(3), 373–409.
- Neta, R. (2022). Credence and belief. *Philosophical Studies*. <https://doi.org/10.1007/s11098-022-01881-8>
- Nozick, R. (1974). *Anarchy, State, and Utopia* (Issue 1, pp. 187–201). Basic Books.
- Olsen, K. (2018). Subjective Rightness and Minimizing Expected Objective Wrongness. *Pacific Philosophical Quarterly*, 99(3), 417–441. <https://doi.org/10.1111/papq.12188>
- Owens, D. (2013). Value and Epistemic Normativity. *Teorema: Revista Internacional de Filosofia*, 32(3), 35–58.
- Parfit, D. (1988). *What We Together Do*. Unpublished Manuscript.
- Pettigrew, R. (2016). Jamesian Epistemology Formalised: An Explication of ‘The Will to Believe’. *Episteme*, 13(3), 253–268.
- Prolini, F. (2019). No Justificatory Closure without Truth. *Australasian Journal of Philosophy*, 97(4), 715–726. <https://doi.org/10.1080/00048402.2018.1564059>
- Radford, C. (1966). Knowledge: By Examples. *Analysis*, 27(1), 1–11. <https://doi.org/10.2307/3326979>
- Raz, J. (2011). *From Normativity to Responsibility*. Oxford University Press.
- Resnik, M. (1987). *Choices: An Introduction to Decision Theory*. University of Minnesota Press.
- Reynolds, S. L. (2013). Justification as the appearance of knowledge. *Philosophical Studies*, 163(2), 367–383.
- Rosenkranz, S. (2021). *Justification as Ignorance: An Essay in Epistemology*. Oxford University Press. <https://doi.org/10.1093/oso/9780198865636.001.0001>
- Ryan, S. (1996). The Epistemic Virtues of Consistency. *Synthese*, 109(2), 121–141. <https://doi.org/10.1007/bf00413765>
- Schoeman, F. (1987). Statistical vs. Direct Evidence. *Noûs*, 21(2), 179–198. <https://doi.org/10.2307/2214913>
- Schroeder, M. (2021). *Reasons First*. Oxford University Press.
- Simion, M., Kelp, C., & Ghijzen, H. (2016). Norms of Belief. *Philosophical Issues*, 26, 374–392.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press UK.
- Smith, M. (2018). When Does Evidence Suffice for Conviction? *Mind*, 127(508), 1193–1218. <https://doi.org/10.1093/mind/fxz026>
- Smith, M. (2022). The Hardest Paradox for Closure. *Erkenntnis*, 87(4), 2003–2028.

- Smithies, D. (2012). The Normative Role of Knowledge. *Noûs*, 46(2), 265–288.
- Sosa, E. (2015). *Judgment and Agency*. Oxford University Press.
- Sosa, E. (2021). *Epistemic Explanations: A Theory of Telic Normativity, and What it Explains*. Oxford University Press.
- Sturgeon, S. (2008). Reason and the grain of belief. *Noûs*, 42(1), 139–165.
- Sutton, J. (2005). Stick to What You Know. *Noûs*, 39(3), 359–396. <https://doi.org/10.1111/j.0029-4624.2005.00506.x>
- Sylvan, K. (2018). Knowledge as a Non-Normative Relation. *Philosophy and Phenomenological Research*, 97(1), 190–222.
- Talbot, B. (2014). Truth Promoting Non-Evidential Reasons for Belief. *Philosophical Studies*, 168(3), 599–618. <https://doi.org/10.1007/s11098-013-0139-1>
- Thomson, J. J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems*, 49(3), 199–219. <https://doi.org/10.2307/1191633>
- Unger, P. K. (1975). Ignorance: A Case for Scepticism (*Issue 105*, p. 369). Oxford University Press.
- Whitcomb, D. (2010). Curiosity was Framed. *Philosophy and Phenomenological Research*, 81(3), 664–687.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Williamson, T. (2007). On Being Justified in One's Head. In M. Timmons, J. Greco, & A. R. Mele (Eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*. Oxford University Press.
- Williamson, T. (2014). Very Improbable Knowing. *Erkenntnis*, 79(5), 971–999.

**How to cite this article:** Dutant, J., & Littlejohn, C. (2023). What is rational belief?. *Noûs*, 1–27. <https://doi.org/10.1111/nous.12456>