

On the Epistemic Rationality and Significance of Self-Fulfilling Beliefs

This is a preprint of an article published in Synthese. The final authenticated version is available online at: <https://doi.org/10.1007/s11229-020-02977-6>

Chad Marxen
Brown University
chad_marxen@brown.edu

Abstract Some propositions are not likely to be true overall, but are likely to be true if you believe them. Appealing to the platitude that belief aims at truth, it has become increasingly popular to defend the view that such propositions are epistemically rational to believe. However, I argue that this view runs into trouble when we consider the connection between what's epistemically rational to believe and what's practically rational to do. I conclude by discussing how rejecting the view bears on three other epistemological issues. First, we're able to uncover a flaw in a common argument for permissivism. Second, we can generate a problem for prominent versions of epistemic consequentialism. Finally, we can better understand the connection between epistemic rationality and truth: epistemic rationality is a guide to true propositions rather than true beliefs.

I. Introduction

Will the sun rise tomorrow? And if you believe it will, will your belief be true?

Normally, the answer is the same to pairs of questions like this one. Your evidence follows suit: it supports the proposition that the sun will rise tomorrow iff it supports the claim that if you were to believe the sun would rise tomorrow, your belief would be true. But in some situations, your evidence supports different answers to these questions.

One kind of situation where this happens is what may be called a “self-fulfilling” situation. As an example, consider a case where you have strong evidence that you will use your toothbrush again—you're still a month away from when you usually purchase a new one. You also have good evidence that if you were to strongly believe that you would never use your toothbrush again, you would throw it in the trash can, in which case you would never use it again.

Now consider the following proposition: you'll never use your toothbrush again. Is this proposition rational for you to believe? Intuitively, it's clear that it isn't. However, you know that if you were to believe it, it would be true. So there's a puzzle: given that belief aims at truth, how could this proposition not be rational for you to believe?¹

In asking whether this (or any) proposition is rational to believe we'll be concerned with *epistemic* rationality, unless noted otherwise. It is the same as what other epistemologists call 'epistemic justification', and it concerns beliefs or credences rather than actions. Further, it is typically, at the very least, tied to evidence. In contrast to practical rationality, it does not depend on preferences.

There's an analogous puzzle for credences. To see this, we can add to the previous case that you have good evidence that if you were to be highly confident that you would never use your toothbrush again, then you would throw it out, and so never use it again. And we can ask an analogous question: is a high credence in the proposition that you'll never use your toothbrush again a rational credence to have? It's clear that it isn't. But again, you know that if you were to have a high credence in the proposition, then it would be true. And since the aim of credence is accuracy, how could a high credence in the proposition not be rational for you to have?

Let's call the content of potentially self-fulfilling credences in these situations, "High-Credence-Fulfillable but Improbable propositions" (HCFI propositions). And let's call the content of potentially self-fulfilling beliefs in these situations, "Belief-Fulfillable but Improbable propositions" (BFI propositions). To be more precise, we'll say that p is a BFI proposition just in case p is improbable overall, but p would be probable if you were to believe it. Similarly, let's say that p is a HCFI proposition just in case p is improbable overall,

¹ For discussions of the thought that belief aims at truth, see Shah (2003), and Velleman (2000).

but p would be probable if you were to have a high credence in it. As to the type of probability in question, I'll always mean *epistemic* probability (i.e. probability on the relevant agent's evidence). By 'probable', I'll always mean *quite* probable (e.g. the probability is higher than .8), and by 'improbable' I'll always mean *quite* improbable (e.g. the probability is lower than .2).

Appealing to the fact that BFI propositions will be fulfilled (i.e. made true) if they're believed (i.e. they're belief-fulfillable), it has become increasingly popular to defend (directly or indirectly) the view that BFI propositions are rational to believe.² And the same goes for the analogous view concerning HCFI propositions. James Joyce provides (and endorses) a typical defense of the view concerning BFI propositions in the following passage:

While this [believing antecedently-improbable propositions] typically indicates an epistemic failing, it need not be problematic when the belief in question causally contributes to its own truth. For example, as Williams James argued, wishful thinking can be legitimate for beliefs that are self-fulfilling prophecies...³

Despite the rising popularity of these views concerning BFI propositions and HCFI propositions, in this paper I will pose a significant problem for them. The problem is a consequence of two claims. First, what's rational to believe is—in a certain sense—

² Raleigh (2015, 2017), and Joyce (2007) argue that belief-fulfillable propositions are rational to believe, even when they're improbable. Dahlback (forthcoming), Drake (2017), Kopec (2015), Peels (2015), Reisner (2007, 2013, 2015), and Velleman (1989) defend the view that belief-fulfillable propositions are rational to believe in cases where their probabilities are unspecified (rather than low). But the principles they employ imply that BFI propositions are rational to believe. Pettigrew (2018), Joyce (2018), and Caie (2013) endorse consequentialist versions of epistemic utility theory that imply that high credences in HCFI propositions are rational to have. Drake (2017), Foley (1991), Reisner (2007, 2013, 2015), Sharadin (2016), and Talbot (2014) all argue that a reason to believe that a proposition would be true if believed, is an *epistemic* reason to believe the given proposition. This view plausibly implies that if you have a great reason of this kind concerning an improbable proposition, then this proposition is epistemically rational for you to believe.

³ See Joyce (2007, p. 558). For a similar defense, see Velleman (1989).

actionable. Second, BFI propositions and HCFI propositions are not actionable in this sense. After arguing for the merits of each claim, I will point out three epistemically significant upshots. First, we're able to uncover a flaw in a common argument for permissivism. Second, we can generate a problem for prominent versions of epistemic consequentialism. Finally, we're able to better understand the connection between epistemic rationality and truth: epistemic rationality is a guide to true propositions rather than true beliefs.

II. Epistemic Rationality's Impact on Practical Rationality

To see why it's plausible that epistemic rationality bears on practical rationality, we'll consider some cases. But before doing so, I should note that when we're wondering whether an action is rational to do we'll be talking about *practical* rationality. While I don't want to take sides on various issues, I do want to commit to a few claims. For starters, practical rationality is paradigmatically concerned with actions rather than beliefs, and it is sensitive to preferences. In contrast to the objective 'ought', practical rationality does not depend on facts beyond one's ken; to employ the distinction between subjective and objective (normative) reasons for action, practical rationality depends on subjective (i.e. perspective-dependent) reasons rather than objective reasons. Now here's the first case:

Vaccination Preparation. A couple of months ago, Soraya and Michael welcomed their first child, Natasha, into the world. As Soraya and Michael care a lot about their daughter, it's rational for Soraya to have Natasha vaccinated, on the condition that it's probably crucial that Natasha be vaccinated. Prior to Natasha's birth, Soraya and Michael attended a course for expecting parents. The course went into great detail concerning the importance of babies receiving vaccines. Soraya and Michael also

discussed this matter with some pediatrician friends, who likewise stressed how vital it is for babies to receive these vaccines.

Question: is it rational for Soraya to have Natasha vaccinated? It certainly seems to be. We know that it's rational for Soraya to have Natasha vaccinated, *on the condition that it's probably crucial that Natasha be vaccinated*. Further, given the strong testimonial evidence that Soraya received, it's rational for her to believe that it is in fact crucial that Natasha be vaccinated. So surely, it must be rational full stop for her to have Natasha vaccinated.

Here's the second case:

Sense of Direction. Rishonda and Evan recently made some new friends and are heading to their place for the first time. Evan is driving today as he usually does. As they're approaching a potentially important exit, they must make a decision whether to take it. Given their plans for the night, it's rational for Evan to take the exit, on the condition that taking the exit is probably fastest. With enough time to take the exit, Rishonda confidently tells him that it's fastest. In Rishonda and Evan's many years together, she's shown herself to be quite reliable in this kind of situation; she has that quintessential sense of direction.

Question: is it rational for Evan to take the exit? It definitely seems to be. We know that it's rational for him to take the exit, on the condition that taking the exit is probably fastest.

Further, we know that it's rational for him to believe that taking the exit *is* fastest. So surely, it must be rational for Evan to take the exit.

Sense of Direction and **Vaccination Preparation** suggest that the following principle is true:

Epistemic Rationality is Actionable (ERIA): if p is rational for you to believe, it's rational for you to act as if p is probably true.⁴

What precisely does ERIA amount to? Let's start with "p is rational for you to believe".

Employing the familiar distinction between *ex ante* and *ex post* epistemic rationality (Goldman 1986), that a particular proposition is rational for you to believe is a matter of *ex ante* epistemic rationality. The propositions that are *ex ante* rational for you to believe may differ from what you *ex post* rationally believe—the propositions that you believe in a rational way. To illustrate, that you're an outgoing person might be rational for you to believe because you have strong evidence that you're outgoing. Nonetheless, you might fail to rationally believe this proposition if you simply don't believe it, or if you only believe it for a bad reason (e.g. that you're a Gemini). Whereas *ex ante* epistemic rationality is focused on potential beliefs, *ex post* epistemic rationality is focused on actual beliefs.

An analogous distinction needs to be made between *ex ante* and *ex post practical* rationality (Whiting 2014). That an action is rational for you to do is a matter of *ex ante* practical rationality. While an action may be rational for you to do, you may not do it in a rational way if you fail to do it, or if you only do it for a bad reason. For example, it may be rational for you to break up with your partner. But if you don't break up with your partner, or if you only break up with your partner for a bad reason (e.g. out of deference to your favorite palm reader), then you didn't rationally break up with your partner—this action isn't *ex post* rational. *Ex post* practical rationality is focused on *actual actions* while *ex ante* practical rationality is centered on *potential actions*.

⁴ For those who endorse views in a similar vein as ERIA, see: Dancy (2000), Fantl and McGrath (2002, 2009), Ganson (2008), Gibbons (2010), Kiesewetter (2011), Lord (2015), Mason (2013), Robertson (2011), Scanlon (2008), and Zimmerman (2008).

Keeping these distinctions in mind, what is it for p to be actionable in the sense that it's rational for you to act as if p is probably true? We'll say that it's rational for you to act as if p is probably true iff the following condition obtains: every action that is rational for you to do, on the condition that p is probably true, is also rational simpliciter for you to do.⁵ So the more cumbersome formulation of ERIA says the following:

ERIA: if p is rational for you to believe, then the following condition holds: every action that's rational for you to do, on the condition that p is probably true, is also rational simpliciter for you to do.

To properly understand ERIA, let's apply it to a case. So suppose it's important that you renew your driver's license. Given your situation, it's rational for you to drive to the Wakefield DMV, on the condition that there's probably a significantly shorter wait time at the Wakefield DMV than the Cranston DMV. Let's also suppose that good testimony makes it rational for you to believe that the wait time at the Wakefield DMV is in fact significantly shorter than at the Cranston DMV. Then surely, it's rational for you to drive to the Wakefield DMV. According to ERIA, this intuitive line of reasoning is legitimate.

While reflecting on cases illuminates ERIA's plausibility, it's also noteworthy that many who have thought about the connection between what's rational to do (or what one ought to do) and your perspective have thought that something like ERIA must be right.⁶ As John Gibbons (2010, p. 336) says, "There must be something to the idea that if it's

⁵ That an action is rational to do, on the condition that a particular proposition is probably true, may differ from whether the action is rational simpliciter to do. It is rational for me to go to the emergency room, on the condition that I probably have appendicitis. But it is not rational simpliciter for me to go to the emergency room.

⁶ See: Dancy (2000), Fantl and McGrath (2002, 2009), Ganson (2008), Gibbons (2010), Kiesewetter (2011), Lord (2015), Mason (2013), Robertson (2011), Scanlon (2008), and Zimmerman (2008). They all think that your epistemic position (rather than your beliefs) determines what you morally (or rationally) ought to do.

reasonable to believe that p then it's reasonable to act as though p." Or as Jeremy Fantl and Matthew McGrath say, "S is justified in believing that p only if S is rational to act as if p."⁷ The intuitiveness of their view can be illuminated through considering the connections between ex ante rationality and reasons. If it's rational for you to believe some claim, you have good reason to believe it. And if you have good reason to believe it, you have good reason to act in ways that are sensible under the assumption that the proposition is probable—you have good reason to act as if the proposition is probable. If you have good reason to act as if it is probable, it's rational for you to act as if it is probable. Putting everything together, if it's rational for you to believe some proposition, it's rational for you to act as if it is probable.

Although many agree that something like ERIA must be true, some such as Fantl and McGrath defend a stronger principle. According to them, if p is rational to believe, then p is *strongly actionable*: every action that's rational to do, on the condition that p is *true*, is also rational simpliciter to do. ERIA only says that if p is rational to believe, then p is *weakly actionable*: every action that's rational to do, on the conditional that p is *probably true*, is also rational simpliciter to do.

Taking into account this distinction is important when considering cases that Jessica Brown and others have used to pose problems for the view that what's rational to believe is strongly actionable.⁸ Consider a modified version of one of Brown's cases where Liz possesses testimony from her parents that makes it rational for her to believe that she was born in England. She's then offered a bet that pays her a dollar if she was born in England,

⁷ See Fantl and McGrath (2002, p. 78). Ganson (2008) also finds this claim intuitively compelling. Fantl and McGrath make it clear that when they use "justified in believing that p," they are talking about ex ante (rather than ex post) justification.

⁸ See Brown (2008). For related counterexamples, see: Reed (2010), and Roeber (2018). Thanks to an anonymous referee for bringing up this worry for ERIA.

while she loses her home if she wasn't. It's rational for her to accept this bet, *on the condition that she was born in England*. But given the chance that she wasn't, it's certainly not rational full stop for her to accept this high-stakes bet. This case involving a high-stakes bet reveals that the proposition that she was born in England is not strongly actionable even though it is rational for her to believe.

For a different kind of case from Brown, consider the previous case except that there's a bet with heavily weighted odds: Liz is offered a bet where she wins a million millionths of a penny if she was born in England, and she loses \$1 if she wasn't. It's rational for Liz to accept this heavily weighted bet, on the condition that the given proposition is true. But given the bet's negative expected utility, it's intuitively not rational full stop for her to accept this bet. So this case involving a heavily weighted bet also reveals that the given proposition is not strongly actionable even though it is rational for her to believe.

A salient feature of these two kinds of cases is that Liz is guaranteed to win each bet, *on the condition that the proposition is true*. So the negative utility of losing is irrelevant to the expected utility of taking the bet *on the condition that the proposition is true*. The bet's conditional expected utility is just equivalent to the positive utility of winning. So it's clearly rational for Liz to accept each bet, on the condition that the given proposition is true. But notice that there's no guarantee that Liz wins each bet, *on the condition that the given proposition is probably true*. So the weighted negative utility of losing a given bet might trump the weighted positive utility of winning it. These cases therefore do not suggest that it's rational for Liz to accept each bet, on the condition that the given proposition is *probably* true. So while these cases suggest that what's rational to believe is not strongly actionable—since the given bets are

clearly not rational full stop to accept—they're compatible with what's rational to believe being weakly actionable.⁹

Brown's case and others make it plausible that practical rationality should not take for granted that what's rational to believe is *true*. This claim is of course independently plausible. After all, many propositions that are rational to believe are false. That said, these cases leave open the possibility that practical rationality takes for granted that what's rational to believe is *probably true*. And it's good that the aforementioned cases leave open this possibility. The reason: **Sense of Direction, Vaccination Preparation**, and many other cases suggest that practical rationality does take for granted that what's rational to believe is probably true.¹⁰

Now in addition to the belief version of ERIA, there's the credence version. According to it, it's rational for you to have a particular credence only if the following condition obtains: every action that's rational for you to do in light of this credence, is also rational simpliciter for you to do. We'll say that A is rational for you to do in light of a credence of C in P just in case the following condition holds: A is rational for you to do, on the condition that P's probability is C. To illustrate, to say that bringing an umbrella on your walk is rational for you to do in light of a credence of .5 in the proposition that there will be a downpour during your walk is to say that bringing an umbrella on your walk is rational for you to do, on the condition that the proposition that there will be a downpour during your walk has a probability of .5.

⁹ While I won't be discussing Roeber's (2018, pp. 177-178) case due to space constraints, the explanation for why his case does not pose a problem for ERIA is the same: the expected utility of the relevant action on the condition that p is true is greater than the expected utility of the action on the condition that p is probably true.

¹⁰ I'm thankful to an anonymous referee for pushing me to clarify why we should think that ERIA holds in general.

To see the plausibility of the credal version of ERIA, let's consider the credal analogue of the aforementioned DMV case. So suppose it's rational for you to have a high credence in the claim that the wait time at the Wakefield DMV is significantly shorter than the Cranston DMV. Further, suppose that driving to the Wakefield DMV is rational for you to do, in light of this credence. Since it's rational for you to have this credence, and the rational thing to do in light of the credence is to drive to the Wakefield DMV, it's plausible that it's rational simpliciter to drive to the Wakefield DMV. According to the credence version of ERIA, this intuitive line of reasoning is warranted.

Further, notice that both versions of ERIA are quite similar in that they describe how *ex ante epistemic* rationality bears on *ex ante practical* rationality; what's rational to do in light of what's rational to think, is also rational simpliciter to do. Given how similar these versions of ERIA are, it's plausible that they should stand or fall together.¹¹

III. Beliefs and Practical Rationality

While ERIA may seem plausible, you might worry about ERIA's implications concerning the relationship between beliefs and practical rationality. According to ERIA, epistemic rationality plays a role in determining practical rationality. Some philosophers have rejected this view, thinking instead that desires and beliefs (or credences) determine what's rational to do.¹² Derek Parfit (2011, p. 112) explains his preferred version of this view in the following passage:

¹¹ The belief version will be used to bear on whether BFI propositions are rational to believe and the credence version will be used to bear on whether high credences in HCFI propositions are rational to have.

¹² See Parfit (2011), and Rinard (forthcoming) for endorsements of this view. Jackson (1991) and Ross (1939) are among those who (in terms of the moral (or deliberative) 'ought') accept analogues of Parfit's view.

Our desires and acts are rational when they causally depend in the right way on beliefs whose truth would give us sufficient reasons to have these desires, and to act in these ways.

For Parfit, what's rational for you to do is the same as what's rational for you to do, given whatever beliefs you have. We can call this view "BDRA" (beliefs determine rational action).

BDRA and ERIA have different verdicts in many cases. While BDRA may seem attractive, there's a serious problem for it. The problem stems from cases where an action makes sense in light of an agent's beliefs, yet the relevant beliefs are deeply irrational. As a real-life example, consider the case of John Hinkley Jr. attempting to assassinate Ronald Reagan.¹³ Hinkley was obsessed with Jodie Foster and was firmly (irrationally) convinced that assassinating Reagan would make Foster love him. Due to his beliefs and desires, he attempted to assassinate Reagan. Intuitively, it was completely irrational for Hinkley to do what he did. That said, it was rational for him to do what he did, assuming his beliefs were true. So BDRA would say it was rational for Hinkley to attempt assassination; but this is highly implausible.

Another way to see the force of the problem for BDRA is by considering a modified version of **Sense of Direction**. Let's add to the case that Evan is a sexist. While he has great reason to believe that Rishonda is quite reliable concerning directions, his sexism leads him to (irrationally) not believe her when she says that it is fastest to take the exit. Instead, he trusts his gut and believes that it is fastest to take another exit.

Question: is it rational simpliciter for Evan to blow past the exit? Intuitively, it isn't. Given Evan's evidence about how likely it is that Rishonda is right, Evan has great reason to think that taking the exit is fastest. And since Evan strongly wants to take the best route to

¹³ For inspiration to use this case, see Arpaly and Schroeder (2014, pp. 72-73).

their destination, blowing past the exit is not a rational thing for him to do. But according to BDRA, Evan's sexist beliefs make blowing past the exit a rational thing for him to do.

Fortunately for ERIA, Rishonda and Evan's case does not pose a problem for it. For ERIA, Evan's sexism doesn't bear on whether it's rational to blow past the exit. ERIA permits agents to do what makes sense in light of what's rational to believe, not what they actually (perhaps irrationally) believe.

Now there's a related worry concerning the credal version of ERIA: ERIA is apparently incompatible with decision theory. After all, decision theory assumes that what's rational to do is determined by your actual credences—regardless of whether they're rational—and actual preferences. But according to ERIA, it's rational to do what makes sense in light of the credences that are rational to have—regardless of whether you actually have them.¹⁴

Notice that the apparent incompatibility between ERIA and decision theory only arises if we take decision theory to be a *complete* theory of practical rationality. Even though decision theory is the most important component of a complete theory of practical rationality, we need not think it is the only component. Decision theory is a theory of rational action, *given certain credences*. While it's extremely important to theorize about practical rationality in this way, there's a further question about which credences should be used when providing a complete theory of practical rationality: should we use the agent's actual credences or should we use the credences that are rational for the agent to have? While some decision theorists assume that we should use an agent's actual credences, decision theory is consistent with either option. So ERIA doesn't imply that decision theory needs to be

¹⁴ I'm grateful to an anonymous referee for bringing up this concern about the relationship between ERIA and decision theory.

revised. Moreover, it's noteworthy that decision theory assumes that an agent's credences are probabilistic. While this assumption fits well with the view that a complete theory of practical rationality should use the credences that are rational for an agent to have, it doesn't fit well with the view that an agent's actual credences should be used. After all, actual credences need not be—and usually aren't—probabilistic.

So ERIA remains plausible. In the next section, we'll consider how ERIA bears on BFI and HCFI propositions.

IV. Practical Rationality and BFI Propositions

Are BFI propositions actionable? To answer this question, let's consider an ordinary situation involving a BFI proposition:

Terrible Workday. Friedrich isn't worried about how his workday will turn out. He knows the chance that today will be terrible is quite low since he meticulously keeps records of how his workdays go. While the chance is low, there have been some terrible workdays for him over the past few years. And almost always, the sole reason why these days are terrible is that he firmly believes that the day will be terrible: whenever he does, this belief makes him depressed, which in turn makes the day terrible. Thankfully, Friedrich has strong reason to think that he will, as usual, avoid the pessimistic belief that today will be terrible.

While his terrible days are rare, they have a particularly damaging impact on him. He always needs to see his therapist the next day during his only free time, his lunch break. Seeing his therapist tomorrow would be quite unfortunate as he has lunch plans with his inflexible in-laws. They would be quite upset if he canceled plans with them. If he did, he'd need to go out of his way to appease them (e.g. buy them a gift certificate for a fancy restaurant), even though he has a tight budget.

Friedrich is aware of the details of his situation. Let T be the proposition that today will be terrible. While his reasons make it quite improbable that T, he knows that if he firmly believes T, then T will be true. In other words, T is a BFI proposition.

Is it rational for Friedrich to act as if it's probable that T? On the condition that T is probably true, it's rational for him to call his therapist to tell her that he needs to meet tomorrow. And since it's rational for him to do this only if it's rational for him to cancel lunch plans with his in-laws, it must be rational for him to cancel these plans. As it's rational for him to cancel these plans, it must be rational for him to buy his inflexible in-laws a gift certificate to a fancy restaurant. On the condition that T is probably true, all of these actions are rational.

But is it rational simpliciter for Friedrich to do these things? This seems hard to believe. As Friedrich's reasons make it quite improbable that T, how could it be rational for him to cancel lunch plans with his in-laws or rational for him to buy them an expensive gift certificate? Doing each of these actions has significant negative consequences for him. It only seems to make sense to bring about these consequences if it is sufficiently probable that T. But given that T is quite improbable, it doesn't make sense for him to bring about these negative consequences. Since it's irrational simpliciter for him to do these things, T isn't actionable.

While **Terrible Workday** provides us with one good case where a BFI proposition isn't actionable, let's consider another case involving a BFI proposition:

Rehab. As a recovering alcoholic, Al has made so much progress that there's only a minute chance he'll relapse within the next few weeks. While Al doesn't currently believe he'll relapse within the next few weeks, he knows that if he firmly believed that he would, he would now start drinking heavily out of sadness. Having done this,

he would have relapsed within the next few weeks. Fortunately for Al, it's quite probable that he won't come to believe that he'll relapse within the next few weeks. It would be quite bad for Al if he were to relapse. For one thing, he would have to take a leave of absence from work to enter a rehab facility. And to get a leave of absence, he would need to ask his boss for it in advance. Additionally, to enter a rehab facility, he would need to ask his family for a lot of money to pay for admission to the facility.

Al knows the details of his situation. Let R be the proposition that Al will relapse within the next few weeks. While it's quite improbable that R, Al knows that if he firmly believes R, then R. In other words, R is a BFI proposition.

Is R actionable, though? On the condition that R is probably true, it's rational for him to ask his boss for a leave of absence and it's rational for him to ask his family for a lot of money to enter a rehab facility. On the condition that R is probably true, these actions are rational for Al to do.

But is it rational simpliciter for Al to do these things? This seems very implausible. If R is quite improbable, how could it be rational for him to ask his boss for a leave of absence and ask his family for a lot of money to enter rehab? Each of these actions has significantly bad consequences for him. It would seem rational to bring about these consequences only if R were sufficiently probable. But since R is quite improbable, it's irrational simpliciter for him to perform these actions. So R isn't actionable. Assuming ERIA, the BFI propositions in **Rehab** and **Terrible Workday** are not rational to believe.

The central lesson of these cases is that being belief-fulfillable does not suffice for being rational to believe. The BFI propositions we've been looking at are not rational to believe, even though they would be true if the agents believed them.

This is the paper's more modest conclusion, but it's plausible that a stronger conclusion also holds: that no BFI proposition is rational to believe. The BFI propositions in **Rehab** and **Terrible Workday** are not actionable precisely because *they're improbable*. But since every BFI proposition is improbable, it's plausible that no BFI proposition is actionable. To see why this is the case, let P be an arbitrary BFI proposition. Plausibly, there's an even odds bet on P where you know that accepting the bet won't make you believe P. In such a case, it's plausible that it's rational for you to accept this even odds bet on P, *on the condition that P is probably true*. After all, the bet has positive conditional (on the condition that P is probably true) expected utility. Nevertheless, it's not rational simpliciter for you to accept this even odds bet on P. The reason: P is quite improbable and accepting the bet won't change this fact.¹⁵ So if we combine ERIA with this stronger conclusion that no BFI proposition is actionable, we get the result that no BFI proposition is rational to believe.

While we've focused on the belief version of ERIA to determine whether there's a BFI proposition that's rational to believe, analogues of **Rehab** and **Terrible Workday** can be used to argue that high credences in HCFI propositions are never rational to have. When we understand actionability as the credence version of ERIA does, we can use analogues of the previous cases to show that these HCFI propositions aren't actionable. So for credences, a similar modest conclusion also holds: being high-credence-fulfillable doesn't suffice for being rational to have a high credence in. And given the analogue of the generalization point concerning BFI propositions, it's plausible that no HCFI propositions are actionable.

¹⁵ Thanks to an anonymous referee for pushing me to clarify the argument from ERIA to the stronger conclusion.

V. Rejecting BFI propositions: Three Epistemic Upshots

Let's suppose we accept the argument for the modest conclusion that being belief-fulfillable doesn't suffice for being rational to believe. If so, there are some epistemically significant upshots.

The first upshot pertains to the question of whether permissivism is true. A number of authors have recently defended an interesting argument for *intrapersonal belief-permissivism*. According to this view, there is more than one doxastic attitude (e.g. belief, disbelief, suspension of belief) that an agent with a single body of evidence is permitted to have towards a particular proposition. Here's Matthew Kopec and Michael Titelbaum's presentation of the argument for this view:

One day, God makes the following proposal: If you believe that it will rain in Canberra tomorrow, then she will make sure it rains in Canberra tomorrow. But if you believe it won't rain tomorrow, then she'll make sure it doesn't...If you believe that it will rain in Canberra, then it certainly will rain, *so that belief is surely justified*. If you believe it won't, it certainly won't, *so that belief is surely also justified*.¹⁶

Let C be the proposition that it will rain in Canberra tomorrow. The argument aims to establish that C is (antecedently) rational for you to believe and that $\sim C$ is also (antecedently) rational for you to believe. You're aware that if you believe C, then C will certainly be true. From this claim, it's supposed to follow that C is justified (rational) for you to believe. But since being belief-fulfillable doesn't suffice for being (antecedently) rational to believe (i.e.

¹⁶ Kopec and Titelbaum (2016, p. 197), my emphasis. For similar arguments for intrapersonal belief-permissivism, see: Dahlbeck (forthcoming), Drake (2017), Kopec (2015), Peels (2015), Raleigh (2015, 2017), Reisner (2007, 2013, 2015), and Velleman (1989).

the modest conclusion is true), it doesn't follow that C is rational for you to believe. So this argument fails to establish intrapersonal belief-permissivism.

A second upshot concerns the question of whether widely endorsed versions of epistemic consequentialism are true.¹⁷ In what follows, I'll argue that some of the most prominent versions of epistemic consequentialism imply that T is rational for Friedrich to believe. But if the argument in this paper is successful, then T is not rational for him to believe. Before explaining which versions of epistemic consequentialism I have in mind, let's begin with Kristoffer Ahlstrom-Vij and Jeffrey Dunn's rough characterization of epistemic consequentialism:

[M]uch of contemporary epistemology has been marked...by a tacit endorsement of *epistemic consequentialism*, the idea that epistemic rightness—denoted by terms such as 'justification' or 'rationality'—is to be understood in terms of conduciveness to epistemic goods.¹⁸

This characterization of epistemic consequentialism leaves it open whether "conduciveness" should be understood in an objective way as reliabilism does, or in a subjective way as epistemic utility theory does.¹⁹ I'd like to focus on subjective versions of epistemic consequentialism, where what's rational to believe is partially determined by your doxastic attitudes, evidence, reasons, knowledge, etc. This characterization is also silent on what the epistemic goods are. I'll follow the most prominent version of subjective epistemic

¹⁷ See Berker (2013) for a discussion of the prominence of epistemic consequentialism among epistemologists.

¹⁸ See Ahlstrom-Vij and Dunn (2018). For sympathizers of epistemic consequentialism, see: Ahlstrom-Vij and Dunn (2014), Caie (2013), Goldman (1986), Greaves (2013), Greaves and Wallace (2006), Joyce (2018), Konek and Levinstein (2019), Oddie (1997), Pettigrew (2018), and Singer (2018, 2019).

¹⁹ For the most developed versions of epistemic utility theory and reliabilism, see Pettigrew (2016) and Goldman (1979), respectively.

consequentialism—veritistic epistemic utility theory—in assuming that the only fundamental epistemic good is accuracy.²⁰ While I'll assume this, what I'll say also applies to versions of epistemic consequentialism that see accuracy as one of the (perhaps many) fundamental epistemic goods. The previous characterization of epistemic consequentialism is also silent on what kind of (doxastic) attitude is under evaluation. In what follows, I mean to include versions of epistemic consequentialism that are about either beliefs or credences. To that end, let's use "accept" to encompass both beliefs and high credences, and let's use "reject" to encompass both disbeliefs and low credences. Further, let's use "Acceptance-Fulfillable but Improbable propositions" (AFI propositions) to refer to improbable propositions that you can expect to fulfill by accepting them.

According to the versions of epistemic consequentialism under consideration, when potential attitudes become adopted, the resultant attitudes will have a certain degree of epistemic goodness (accuracy). Just as the resultant attitudes vary in their epistemic goodness, our expectations concerning their resultant epistemic goodness also vary. For these versions of epistemic consequentialism, an attitude is rational for you to adopt just in case it is unsurpassed in terms of your expectation of its resultant epistemic goodness (accuracy).²¹

Having clarified which versions of epistemic consequentialism are under consideration, the important point is that these versions plausibly imply that AFI propositions are rational to accept. We can see the plausibility of this claim by revisiting **Terrible Workday**. To see how epistemic consequentialism bears on this case, let's make

²⁰ Those sympathetic with veritism include the following: Alston (2005), Bonjour (1985), Caie (2013), Foley (1993), Goldman (1999), Greaves (2013), Greaves and Wallace (2006), Konek and Levinstein (2019), and Pettigrew (2016, 2018).

²¹ For a similar (informal) characterization of epistemic consequentialism, see Greaves (2013).

the probabilities in it more precise. In particular, let's suppose that the (unconditional) probability of T is .05. Let's also assume that T's probability is .99, on the condition that Friedrich comes to *accept* T. Finally, let's suppose that \sim T's probability is .96, on the condition that he comes to *reject* T.

According to the aforementioned versions of epistemic consequentialism, a particular attitude towards T is rational for Friedrich to adopt just in case this attitude is unsurpassed in terms of his expectation of its resultant accuracy. To determine what attitude towards T is rational for Friedrich to adopt, there are three possibilities to consider: accepting T, a middling attitude towards T, and rejecting T. Let's consider each in turn. Friedrich's expectation of the resultant accuracy of accepting T is quite high; he's aware that T's probability is .99, on the condition that he comes to accept T. But Friedrich's expectation of the resultant accuracy of a middling attitude towards T is not high; the resultant accuracy of this middling attitude is guaranteed to be similarly middling. Finally, Friedrich's expectation of the resultant accuracy of rejecting T is quite high; he's aware that \sim T's probability is .96, on the condition that he comes to reject T. So Friedrich's expectation of the resultant accuracy of rejecting T is easily higher than his expectation of a middling attitude towards T. Nonetheless, his expectation of the resultant accuracy of rejecting T is lower than his expectation of accepting T. Notice that T's probability is .99, on the condition that he comes to accept T. But \sim T's probability is only .96, on the condition that he comes to reject T. So accepting T is unsurpassed in terms of Friedrich's expectations. Therefore, according to the aforementioned versions of epistemic consequentialism, T is rational for Friedrich to accept.²²

²² Greaves (2013), Joyce (2018), Pettigrew (2018), and Raleigh (2017) embrace this verdict (in similar cases) due to their sympathies with epistemic consequentialism.

But as we've already seen, T is not actionable. On the condition that T is probable, it's rational for Friedrich to cancel lunch plans with his in-laws. But given how improbable T is, it is not rational simpliciter for Friedrich to do this; doing so will bring him gratuitous harm. So T is not actionable. Therefore, contrary to the aforementioned versions of epistemic consequentialism, T is not rational for Friedrich to accept. So there's a serious problem for some—but not all—versions of epistemic consequentialism.²³

A third upshot concerns the relationship between epistemic rationality or justification, and truth. There's a common belief that epistemic rationality or justification is a guide to truth. Laurence BonJour articulates this belief in the following passage:

What makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavors is *truth*...bringing it about that our beliefs are epistemically justified will also tend to bring it about...that they are true. If epistemic justification were not conducive to truth in this way, if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth. It is only if we have some reason for thinking that epistemic justification constitutes a path to truth that we as cognitive beings have any motive for preferring epistemically justified beliefs to epistemically unjustified ones.²⁴

According to BonJour, epistemic justification or rationality is a guide to truth. It's unclear, though, how exactly we should understand the claim that epistemic rationality is a guide to

²³ The modest conclusion tells against versions of epistemic consequentialism endorsed by Caie (2013), Greaves (2013), Greaves and Wallace (2006), Joyce (2018), and Pettigrew (2018). That said, the modest conclusion doesn't tell against versions of epistemic consequentialism that Konek and Levinstein (2019) and Carr (2017) are sympathetic with.

²⁴ See Bonjour (1985, pp. 7-8). For similar thoughts see: Cohen (1984), Conee (1992), Horowitz (2019), and Schoenfield (2015).

truth. One way to understand the claim is that epistemic rationality guides us toward the propositions that are likely to be true (independent of whether we believe them)—epistemic rationality is a guide to *true propositions*. Another way to understand the claim is that epistemic rationality guides us toward the propositions that are likely to be true *if believed*—epistemic rationality is a guide to *true beliefs*.

Let's call the view that epistemic rationality is a guide to true propositions and that being (likely to be) a true proposition is a fundamental explanatory property of epistemic rationality, *truth-aimed epistemology*. And let's call the view that epistemic rationality is a guide to true beliefs and that being (likely to be) a true proposition if believed is a fundamental explanatory property of epistemic rationality, *true-belief-aimed epistemology*.

Truth-aimed epistemology is suggested by the modest conclusion of this paper that being belief-fulfillable doesn't suffice for being rational to believe. If true-belief-aimed epistemology were correct, then it is hard to see why being belief-fulfillable would not suffice for being rational to believe.

Truth-aimed epistemology also fits nicely with the thought that epistemic rationality is about conforming your mind (doxastic attitudes) to the way the world is. Epistemic rationality begins with what's likely to be true and concludes with verdicts about what's rational to believe. Whereas practical rationality is about making the world conform to your mind (desires), epistemic rationality is about making your mind (doxastic states) conform to the world.²⁵

Finally, truth-aimed epistemology can provide a solution to the puzzle mentioned at the beginning of the paper. The puzzle arises in cases where some proposition is quite

²⁵ For further discussion about using the direction of fit metaphor to distinguish practical rationality from epistemic rationality, see Konek and Levinstein (2019).

improbable, but you know that if you were to believe the proposition it would be true. The puzzle is this: given that belief aims at truth, how could it not be rational for you to believe the proposition? The key to solving the puzzle has to do with how we understand “belief aims at truth.”²⁶ If we understand the expression as an implicit endorsement of true-belief-aimed epistemology, then there’s no clear answer to the question. But if we understand the expression as an implicit endorsement of truth-aimed epistemology, then we can answer the question. What’s likely to be true is explanatorily fundamental for epistemic rationality rather than what’s likely to be true if believed. So the fact that a proposition is likely to be true if believed is epistemically irrelevant. Epistemic rationality is a matter of conforming your mind to the world rather than using your mind as an instrument to produce a conformity between it and the world.²⁷

VI. Conclusion

BFI propositions are philosophically significant because they generate the aforementioned puzzle. They’re also significant because potential solutions to this puzzle bear on other philosophical issues such as permissivism, epistemic consequentialism, and the connection between epistemic rationality and truth. Given the intrinsic philosophical significance of BFI propositions, along with their connection to these other philosophical issues, it’s worth thinking carefully about them.

As I’ve argued, we should resist the tempting argument for the conclusion that BFI propositions are rational to believe. Instead, we should think that BFI propositions aren’t

²⁶ Antill (2020) also thinks that the expression is ambiguous; it can be understood as saying that the aim of belief is acquiring true beliefs or it can be understood as saying that the aim of belief is believing what is true. With this ambiguity in mind, he argues that Velleman’s (1989) argument for (epistemic) permissivism is unsuccessful.

²⁷ Even according to truth-aimed epistemology, of course, epistemic rationality will typically guide you to true beliefs. But at a fundamental level, epistemic rationality aims at the (belief-independent) truth.

rational to believe. The argument for this conclusion hinged on two claims: what's rational to believe is actionable, and BFI propositions are not actionable.

Finally, if BFI propositions are not rational to believe, truth-aimed epistemology provides a reasonable explanation of this fact. I've argued that BFI propositions provide us with a reason to accept truth-aimed epistemology over true-belief-aimed epistemology. They suggest that epistemic rationality is most fundamentally a guide to true propositions, not a guide to acquiring true beliefs.

Acknowledgements Many thanks to Nomy Arpaly, Zach Barnett, Jamie Dreier, Louis Gularte, Richard Pettigrew, Wes Siscoe, and an anonymous referee for *Synthese*. I'm especially grateful to David Christensen and Josh Schechter for their many helpful suggestions.

References

- Alston, W. (2005). *Beyond "Justification": Dimensions of epistemic evaluation*. Ithaca: Cornell University Press.
- Antill, G. (2020). Epistemic freedom revisited. *Synthese*, 197(2), 793-815.
- Arpaly, N., & Schroeder, T. (2014). *In praise of desire*. Oxford University Press.
- Ahlstrom-Vij, K., & Dunn, J. (2014). A defence of epistemic consequentialism. *The Philosophical Quarterly*, 64(257), 541-551.
- Ahlstrom-Vij, K., & Dunn, J. (2018). Introduction. In J. Dunn & K. Ahlstrom-Vij (Eds.), *Epistemic Consequentialism* (pp. 1-22). Oxford: Oxford University Press.
- Berker, S. (2013). Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122(3), 337-393.
- BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42, 167-189.
- Caie, M. (2013). Rational probabilistic incoherence. *Philosophical Review*, 122(4), 527-575.
- Carr, J. (2017). Epistemic utility theory and the aim of belief. *Philosophy and Phenomenological Research*, 95(3), 511-534.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46, 279-295.
- Conee, E. (1992). The truth connection. *Philosophy and Phenomenological Research*, 52(3), 657-669.
- Dahlback, M. (forthcoming). Infinitely permissive. *Erkenntnis*.
- Drake, J. (2017). Doxastic permissiveness and the promise of truth. *Synthese*, 194, 4897-4912.
- Dancy, J. (2000). *Practical reality*. Oxford University Press.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics and justification. *Philosophical Review*, 111(1), 67-94.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Foley, R. (1991). Evidence and reasons for belief. *Analysis*, 51(2), 98-102.
- Foley, R. (1993). *Working without a net: A study of egocentric epistemology*. New York: Oxford University Press.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139, 441-458.

- Goldman, A. (1979). What is justified belief? In G.S. Pappas (Ed.), *Justification and knowledge* (pp. 1-23). Boston: D. Reidel.
- Goldman, A. (1986). *Epistemology and cognition*. Harvard University Press.
- Goldman, A. (1999). *Knowledge in a social world*. Oxford University Press.
- Greaves, H. (2013). Epistemic decision theory. *Mind*, 122(488), 915–952.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459), 607–632.
- Gibbons, J. (2010). Things that make things reasonable. *Philosophy and Phenomenological Research*, 81(2), 335–61.
- Horowitz, S. (2019). The truth problem for permissivism. *The Journal of Philosophy*, 116(5), 237-262.
- Joyce, J.M. (2007). Are Newcomb problems really decisions? *Synthese*, 156(3), 537-562.
- Joyce, J.M. (2018). Accuracy, ratification, and the scope of epistemic consequentialism. In J. Dunn & K. Ahlstrom-Vij (Eds.), *Epistemic consequentialism* (pp. 240-268). Oxford: Oxford University Press.
- Kiesewetter, B. (2011). “Ought” and the perspective of the agent. *Journal of Ethics and Social Philosophy*, 5(3), 1-24.
- Konek, J., & Levinstein, B. (2019). The foundations of epistemic decision theory. *Mind*, 128(509), 69-107.
- Kopec, M. (2015). A counterexample to the uniqueness thesis. *Philosophia*, 43, 403-409.
- Kopec, M., & Titelbaum, M.G. (2016). The uniqueness thesis. *Philosophy Compass*, 11, 189-200.
- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101(3), 461-482.
- Lord, E. (2015). Acting for the right reasons, abilities, and obligation. *Oxford Studies in Metaethics*, 10, 26–52.
- Mason, E. (2013). Objectivism and prospectivism about rightness. *Journal of Ethics and Social Philosophy*, 7(2), 1–21.
- Oddie, G. (1997). Conditionalization, cogency, and cognitive value. *The British Journal for the Philosophy of Science*, 48(4), 533-541.
- Parfit, D. (2011). *On what matters*. Oxford University Press.
- Peels, R. (2015). Believing at will is possible. *Australasian Journal of Philosophy*, 93(3), 524-541.
- Pettigrew, R. (2016). *Accuracy and the laws of credence*. Oxford University Press.

- Pettigrew, R. (2018). Making things right: The true consequences of decision theory in epistemology. In J. Dunn & K. Ahlstrom-Vij (Eds.), *Epistemic consequentialism* (pp. 220-239). Oxford: Oxford University Press.
- Raleigh, T. (2015). An argument for permissivism from safespots. In W. van der Hoek, W.H. Holliday, & W-f Wang (Eds.), *Logic, rationality, and interaction* (pp. 308-315). Dordrecht: Springer.
- Raleigh, T. (2017). Another argument against uniqueness. *The Philosophical Quarterly*, 67(267), 327-346.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44(2), 224-244.
- Reisner, A. (2007). Evidentialism and the numbers game. *Theoria*, 73(4), 304-316.
- Reisner, A. (2013). Leaps of knowledge. In T. Chan (Ed.), *The aim of belief* (pp. 167-183). Oxford University Press.
- Reisner, A. (2015). A short refutation of strict normative evidentialism. *Inquiry*, 58(5), 477-485.
- Rinard, S. (forthcoming). Pragmatic skepticism. *Philosophy and Phenomenological Research*.
- Robertson, S. (2011). Epistemic constraints on practical normativity. *Synthese*, 181(1), 81–106.
- Roeber, B. (2018). The pragmatic encroachment debate. *Noûs*, 51(1), 171-195.
- Scanlon, T. M. (2008). *Moral dimensions: Permissibility, meaning, blame*. Cambridge, MA: Belknap Press.
- Schoenfield, M. (2015). Bridging rationality and accuracy. *The Journal of Philosophy*, 112(2), 633-657.
- Shah, N. (2003). How truth governs belief. *Philosophical Review*, 112(4), 447-482.
- Sharadin, N.P. (2016). Nothing but the evidential considerations? *Australasian Journal of Philosophy*, 94(2), 343-361.
- Singer, D.J. (2018). How to be an epistemic consequentialist. *The Philosophical Quarterly*, 68(272), 580-602.
- Singer, D.J. (2019). Permissible epistemic trade-offs. *Australasian Journal of Philosophy*, 97(2), 281-293.
- Talbot, B. (2014). Truth promoting non-evidential reasons for belief. *Philosophical Studies*, 168, 599-618.
- Ross, W.D. (1939). *Foundations of ethics*. Oxford University Press.
- Velleman, D. (1989). *Practical reflection*. Princeton University Press.
- Velleman, D. (2000). *The possibility of practical reason*. Oxford University Press.
- Whiting, D. (2014). Keep things in perspective: Reasons, rationality and the a priori. *Journal of Ethics and Social Philosophy*, 8(1), 1-22.

Zimmerman, M.J. (2008). *Living with uncertainty: The moral significance of ignorance*. Cambridge University Press.