

Epistemic Utility Theory's Difficult Future

This is a preprint of an article published in Synthese. The final authenticated version is available online at: <https://doi.org/10.1007/s11229-021-03120-9>

Chad Marxen
Brown University
chad_marxen@brown.edu

Abstract According to epistemic utility theory, epistemic rationality is teleological: epistemic norms are instrumental norms that have the aim of acquiring accuracy. What's definitive of these norms is that they can be expected to lead to the acquisition of accuracy when followed. While there's much to be said in favor of this approach, it turns out that it faces a couple of worrisome extensional problems involving the future. The first problem involves credences about the future, and the second problem involves future credences. Examining prominent solutions to a different extensional problem for this approach reinforces the severity of the two problems involving the future. Reflecting on these problems reveals the source: the teleological assumption that epistemic rationality aims at acquiring accuracy.

I. Introduction

It's natural to assume that there's a close connection between epistemic rationality and truth.¹ This natural assumption applies to credences as well as beliefs. Even though credences can't be true or false, they can be more or less accurate.

Assuming there's a close connection between the epistemic rationality of a credence and its accuracy, it would be valuable to have a theory that explains what the connection consists in. The standard version of epistemic utility theory (EUT) endorses Credal Veritism—the claim that there's exactly one fundamental epistemic good: accuracy. Advocates of the standard version also usually hold that the rational credal state is the one that is most preferable in terms of expected epistemic goodness, if you adopt it.² If we combine these views, we get the result that the rational credal state is the one that is most preferable in terms of expected accuracy, if you adopt it.

¹ Hereafter, by 'rational', I mean epistemically rational (instead of, for example, practically rational).

² This view isn't unanimously held, though. See Konek & Levinstein (2019) for a version of EUT where credences are not evaluated on the condition that they're adopted.

For this version of EUT, epistemic rationality is sensitive to your conditional expectation concerning how accurate a credal state is—your expectation of the accuracy of a credal state, if you adopt it.³ This is in keeping with EUT’s teleological assumption that epistemic rationality aims at acquiring accuracy.

In this paper I focus on this conditional sort of EUT (CEUT) motivated by its teleological assumption. We’ll begin with a brief introduction to CEUT and then I’ll pose a difficult problem for it based on cases involving self-fulfilling credences about the future. A salient aspect of the problem is that self-fulfilling credences influence their accuracy. This is significant because plausible solutions have been given to a similar and challenging problem for CEUT: Hilary Greaves’ *Trade-Off Problem* (2013). Since this problem is also focused on cases where credences influence their accuracy, it seems apt to explore whether plausible solutions to it will also help with our problem. As we’ll come to see that these solutions won’t help, we’ll be able to appreciate the force of our problem. I’ll then present a second problem for CEUT involving future credences. It’ll also turn out that the promising solutions to the Trade-Off Problem won’t help with our second problem. I’ll conclude by diagnosing the source of these problems for CEUT: the conditional aspect of it, which flows from its teleological assumption that epistemic rationality aims at acquiring accuracy.

II. The Problem of Self-Fulfilling Credences

To have a fully fleshed-out theory, CEUT needs to answer two questions: (1) What does EU consist in? (2) And what instrumental norms govern an agent’s pursuit of it?

According to CEUT, EU is just accuracy (Credal Veritism). A natural follow-up question for credal veritists is how credal accuracy should be measured. The most popular

³ See Pettigrew (2016) and (2018) for the most developed version of this conditional kind of EUT. For others sympathetic with conditional EUT, see: Caie (2013), Greaves (2013), and Joyce (2018).

approach uses the Brier score, which positively evaluates a credence in a truth insofar as the credence approaches certainty, and positively evaluates a credence in a falsehood insofar the credence approaches complete doubt.⁴

What instrumental norms govern an agent's pursuit of EU? While standard versions of CEUT endorse many norms, we'll focus on the maximization norm: if it is rational to occupy credal state C, then occupying C maximizes expected epistemic utility (EEU) (by the lights of C).⁵ The EEU of a credal state A, by the lights of C, is equivalent to the weighted sum of the EU scores of A at each world, where the EU score of A at a given world is weighted by the probability that C assigns to the given world, given that A is adopted.⁶

With CEUT's maximization norm in mind, let's see how it bears on the following case:

Painful Prep. Sony is required to take an Introduction to Proofs course in the fall. Unfortunately, the course requires some highly unpleasant prep work. Sony knows that if he does the prep work, it will be smooth sailing from there on out. But if he doesn't, the class will be a nightmare. His history indicates that he will do the prep

⁴ More precisely, the inaccuracy of a credence at a world is equal to the squared distance between the credence and the maximally accurate credence at the world. The shorter the distance, the less inaccurate the credence is; the longer the distance, the more inaccurate it is. Multiplying the inaccuracy score by -1 results in an accuracy score for the given credence. The closer the score is to 0, the more accurate the credence is, and the farther it is from 0, the less accurate it is. A perfectly accurate credence in a true proposition is the maximal credence in it. A perfectly inaccurate credence in a true proposition is the minimal credence in it. The overall accuracy for a set of credences (i.e. a credal state) is the average of the accuracy scores for the given credences.

⁵ For a discussion of these additional norms, see Pettigrew (2016).

⁶ Those who accept a causal maximization norm think that the principle should take into account the probability of worlds if the agent were to occupy the given credal state. Those who accept an evidential maximization norm think that the principle should take into account the probability of worlds if the agent in fact comes to occupy the given credal state. While there's this difference between causal and evidential CEUT, they agree in thinking that an agent's occupation of a credal state should bear on the probability of worlds. For the rest of the paper I will be focusing on causal CEUT, but the success of the arguments below does not hinge on which version of CEUT is correct. It's worth noting that the leading defenders of CEUT (e.g. Pettigrew (2018), and Joyce (2018)) subscribe to causal versions of it.

work, as long as he doesn't become too pessimistic about the course. But he also knows that things won't go so well if he does become pessimistic: whenever he's highly confident that something will go badly, he's unable to muster the will to do unpleasant preparations for it. Unfortunately, Sony has been in this difficult situation many times; his history clearly indicates that there's a 50/50 chance that he'll manage to avoid this crippling pessimism.

What credence is it *intuitively* rational for Sony to have in the claim that (B): Sony will do badly in Introduction to Proofs? He knows that whether he will do badly is settled by whether he will complete the unpleasant prep work. Since he has strong reason to think there's a 50/50 chance he will complete it, it seems rational for Sony to have a .5 credence in B.

But what credence does CEUT recommend? To answer this, let's first determine the EEU of the intuitively rational credence of .5. There are two (kinds of) worlds to consider: the worlds where B is true and the worlds where it is false. For each world, the .5 credence will have -.25 EU.⁷ Therefore, the EEU of the .5 credence in B is -.25.

What's the EEU of Sony's having a high credence in B? As an illustrative example, we'll calculate the EEU of Sony's having a credence of 1 in B, since the choice of a particular high credence in B is unnecessary to reveal CEUT's implication in Sony's case. At the world where B is true, the credence is perfectly accurate, so the EU of the credence at this world is 0. At the world where B is false, the credence is perfectly inaccurate, so the EU of the credence is -1. The next crucial pieces of information we need are the conditional probabilities of each world. According to CEUT, we are to consider the probabilities of these worlds given that Sony adopts credence 1 in B. But Sony knows that if he's highly

⁷ The accuracy score at each world: $-(1 - (1/2))^2$ or $-(0 - (1/2))^2$, according to the Brier Score.

confident in B, it's extremely likely that he won't do the prep work, and he will do badly in the course. So if he adopts credence 1 in B, it's extremely likely that B will be true. Thus the probability of the world where this credence is inaccurate is extremely low, and the EEU for this credence will be very close to 0—clearly better than the -.25 score we obtained for the intuitively rational credence of .5. So CEUT has the troubling implication that it is epistemically irrational for Sony to adopt the intuitively rational credence of .5.⁸

III. The Severity of the Problem

⁸ Carr (2017) also presents a case involving self-fulfilling credences to argue that CEUT clashes with Bayesianism's central diachronic norm: Conditionalization. While we won't be relying on the opposite kind of case (involving self-frustrating credences), she uses it—employing Caie's (2013) argument—to show that CEUT clashes with Bayesianism's central synchronic norm: Probabilism. Now although we both use self-fulfilling cases to arrive at similar conclusions, we rely on different premises and there's an important evidential difference between our cases. To illustrate the first difference, while Carr leverages our beliefs about the epistemic rationality of credal-state transitions via Conditionalization, we rely on our intuitive judgments about the epistemic rationality of credal states at a single time. Shifting to the second difference, in Carr's case it's not obvious what credence in p the evidence points toward when the agent learns that her future credence in p will match p 's future objective chance. The situation of Sony in **Painful Prep** differs in that his evidence obviously points toward a credence of .5 in B. Finally, this paper attempts to illuminate just how unintuitive CEUT's verdicts are in a plethora of ordinary situations.

Greaves (2013) uses a case involving self-fulfilling credences too. However, in her case, the agent's evidence doesn't point towards a particular credence in the relevant proposition. Instead, the agent knows that whatever credence she comes to have in it, her credence will be in accord with the objective chance that it is true. This is noteworthy since if we're dealing with this (or Carr's) case where there's no fact of the matter concerning (or it's unclear) which credence is supported by the evidence, we don't have a strong intuitive reaction (about epistemic rationality) to the case. So, this type of case won't pose an intuitive problem for CEUT. But if we're dealing with our type of self-fulfilling case where there's a fact of the matter (and it's clear) which credence is supported by the evidence, we do have a strong intuitive reaction to it. These cases can be used to—as we do in this paper—generate an intuitive problem for CEUT.

Berker (2013a, 2013b) similarly uses a self-fulfilling example to target epistemological accounts that resemble reliabilism—where beliefs are justified if they're formed in a truth-conducive way, even if the agent has no idea that it is truth-conducive. So the patient who believes against her evidence that she will recover, in a situation where (unbeknownst to her) that belief will make her recovery likely, is justified. This is clearly counterintuitive—but it doesn't apply to cases where the agent has evidence to think that she will recover if she believes she will. Berker (2013b) says, in a footnote, that he believes one could construct troubling cases where the agent is aware of the self-fulfilling nature of believing in her recovery, but adds that arguing for this “takes a bit of care,” and doesn't pursue it. However, from the perspective of CEUT, it is precisely the (rational) expectation of self-fulfilling credal states that is crucial to epistemic rationality.

The Trade-Off Problem and our problem both rely on cases where credences influence their accuracy. Since promising solutions have been given to the Trade-Off Problem—CEUT’s most pressing extensional problem so far—it seems appropriate to see whether these solutions will help with our problem. Additionally, in considering whether any of these promising solutions will work, we can assess our problem’s severity. Now here’s a modified version of the case that generates the Trade-Off Problem (Greaves 2013):

Imps. Emily is walking through the Garden of Epistemic Imps when she sees a child playing on the grass in front of her. Emily knows that in a nearby house are five children, each of whom may or may not come out to play in a minute. They can read Emily’s mind, and their algorithm for deciding whether to play outdoors is as follows: if Emily has a credence of 0 in the claim that there is now a child before her, they will surely come out to play. If she has any other credence in the claim that there is a child before her, they will each flip a fair coin to determine whether they will come out to play. Emily must decide which credence to adopt concerning the proposition that there is now a child before her, as well as the five propositions concerning whether each child will be outdoors in a few minutes.

Emily must choose whether to follow her evidence concerning whether there is now a child before her. If she does and is certain that there is now a child before her, then she must also embrace a middling .5 credence in each of the five propositions concerning whether a particular summerhouse child will come out to play. If Emily follows her evidence, she’s guaranteed to have one perfectly accurate credence and five credences that are -.25 accurate, which would mean that her overall accuracy would be approximately -.21.

Emily’s evidence also entails that if she were to reject her evidence concerning whether there is now a child before her, embracing a credence of 0, then each of the five

summerhouse children would surely come out to play. Emily could have one perfectly inaccurate credence along with five perfectly accurate credences. If she adopted this credal state, her overall accuracy would be approximately .17. Since a credal state is rational only if it maximizes EEU, CEUT says that it is irrational for Emily to adopt the evidence-respecting credal state. But as it seems perfectly rational for Emily to believe what she sees, this verdict of CEUT is intuitively incorrect.⁹

Let's begin by considering the *Unconditional Solution* to the Trade-Off Problem (Konek & Levinstein, 2019). While we could assess the expected accuracy of a credal state C using an agent's unconditional credences, CEUT uses the agent's conditional credences, conditional on the agent coming to adopt C. The significance of this distinction is revealed in a case like **Imps** where the expected accuracy of the evidence-disrespecting credal state is only better than the evidence-respecting credal state when using Emily's *conditional* credences instead of her *unconditional* credences.¹⁰ Accordingly, it's worth questioning CEUT's assumption that we should use conditional credences when evaluating credal states.

To this end, the Unconditional Solution appeals to Anscombe's direction-of-fit metaphor to distinguish epistemic rationality from practical rationality in the following way: while epistemic rationality directs us to conform our mind (doxastic states) to the world, practical rationality directs us to conform the world to our mind (desires).¹¹ Assuming this characterization of the distinction between epistemic and practical rationality is correct,

⁹ **Imps** is reminiscent of similar trade-off cases used to argue against epistemic consequentialism as applied to beliefs rather than credences. In these cases the agents can sacrifice the EU of a particular belief in order to acquire a collection of beliefs that has an unsurpassed amount of EU. **Imps** is an instance of this general problem for epistemic consequentialism. For more on trade-off cases, see: Berker (2013a), (2013b), Caie (2013), Firth (1981), and Jenkins (2007).

¹⁰ The evidence-disrespecting credal state in **Imps** is the one consisting of a credence of 0 in the proposition that there is now a child before her and a credence of 1 in the other five propositions.

¹¹ For more on the direction-of-fit metaphor, see Anscombe (1957).

epistemic rationality does not tell us to use our minds as instruments to produce a conformity between them and the world. Accordingly, we should not assess the expected accuracy of a credal state C on the condition that C is adopted. Instead, we should assess the expected accuracy of C using the agent's unconditional credences. We can call this theory 'unconditional' EUT.

While unconditional EUT allows advocates of EUT to hold on to EUT in the face of the Trade-Off Problem, unconditional EUT faces its own problem. The problem arises due to EUT's foundational assumption which Jason Konek and Ben Levinstein state as follows (2019, p. 69, my emphasis): "Norms of epistemic rationality...have their binding force in virtue of the following fact: *they are good means toward the end of securing accuracy.*" Or as Greaves and David Wallace say (2006, p. 610, my emphasis):

According to the cognitive decision-theoretic approach, epistemic rationality consists in taking steps that can reasonably be expected *to bring about epistemically good outcomes* [i.e. outcomes involving accurate credal states].

Following these advocates of EUT, EUT assumes a *teleological* theory of epistemic rationality where epistemic norms are *instrumental* norms that have the epistemic aim of acquiring accuracy. According to EUT, what's definitive of epistemic norms is that you can expect to best satisfy epistemic rationality's aim (i.e. acquiring accuracy) if you follow these norms.

The problem for unconditional EUT is that it's plausible that EUT's teleological foundation can't be reconciled with the claim credal states should be evaluated using unconditional credences (rather than conditional credences).¹² Assuming epistemic norms help you best acquire accuracy, then surely the epistemic norms take into account how your

¹² For similar problems concerning unconditional EUT, see: Carr (2017), Pettigrew (2018), and Talbot (2014).

coming to occupy a credal state will positively (or negatively) bear on how accurate the given credal state is. After all, if epistemic rationality truly aims at *acquiring accuracy*, how could the norms ignore this information?

It's important to note that advocates of unconditional EUT have responded to this problem. They concede that unconditional EUT will prohibit agents in self-fulfilling cases from adopting credal states that the agents expect to do best in terms of acquiring accuracy. That said, advocates of unconditional EUT think they can explain why epistemic norms should ignore information about how your coming to occupy a credal state bears on how accurate it will be, in the following way (Konek & Levinstein, 2019, p. 104):

The fact that [unconditional EUT] sometimes requires rational agents to leave accuracy on the table does not show that its recommended epistemic policies are not a good means to the end of accuracy...The reason: credal states are simply not epistemically valuable in virtue of casually influencing the world, so as to make themselves accurate.

The thought is that since credal states are not epistemically valuable in virtue of their causal properties (they're only epistemically valuable in virtue of being accurate), the epistemic norms governing pursuit of epistemic value should ignore this causal information.

But it's quite unclear why we should infer the conclusion from the premise. To illustrate the worry with the inference, consider a case where you're a car trader and you're deciding between acquiring two cars. You're aware that while the cars have the same market value (prior to acquiring them), one is used while the other is new. So you know that the market value is higher for the used car than the new car, on the condition that you acquire them. But following the Unconditional Solution, a car is not valuable in virtue of how it is caused to have its market value; it's valuable in virtue of its market value. Nonetheless,

practical norms governing your pursuit of a valuable car should clearly take into account how your acquisition of a car bears on its value.

The general lesson here is that a particular (e.g. causal) property being a source of some value (e.g. epistemic, practical, aesthetic) is one matter, and the given property bearing on how the given value should be pursued is an entirely different matter. Accordingly, even if some property is not a source of some value, it does not follow that the property does not bear on how the given value should be pursued. In sum, it's plausible that the Unconditional Solution can't be reconciled with EUT's foundational assumption that epistemic rationality is teleological. So, the search for a promising solution to the Trade-Off Problem continues.

IV. The Error-Theory Solution

According to Richard Pettigrew's *Error-Theory Solution*, we find it counterintuitive that Emily cannot follow her evidence because of the 'evidentialist' heuristic we use. To employ this heuristic, you check to see whether a person's credal state fits her evidence, and conclude that the credal state is rational for her to adopt just in case it fits her evidence. The evidentialist heuristic is a lot easier to use than EUT's accuracy-based rules, which require completing the difficult task of computing the expected accuracy of credal states. But while there's a big benefit in employing the evidentialist heuristic, it comes with a cost. Although it is extremely reliable, it delivers incorrect verdicts in bizarre cases like **Imps** where a person's credal state bears on its accuracy. In the non-bizarre cases where an agent's credal state does not bear on its accuracy, the verdicts of the evidentialist heuristic align with the true verdicts of CEUT. A key assumption of this response is that the cases where you'd expect a potential credal state to bear on its accuracy are rare since they are bizarre; since the cases are so rare, the costs over the long run are small.

Unlike the Unconditional Solution, the Error-Theory Solution looks quite promising. When we're considering **Imps**, we're supposed to trust our judgments concerning a bizarre case. Pettigrew, Konek and Levinstein all describe it as "pathological", while Joyce characterizes it as "fantastical."¹³ Given the strangeness of the case, it's fair to wonder whether we should trust our intuitions in this kind of case. So if our intuitions only go awry in a few bizarre cases and we can explain why this is the case, it looks like the Error-Theory Solution is a promising solution to the Trade-Off Problem.

But can the Error-Theory Solution handle the Problem of Self-Fulfilling Credences? A first worry with applying this solution to our problem is that it's significantly harder to be skeptical concerning our intuitions in a mundane case like **Painful Prep**. Another potential worry with applying this solution to our problem turns on how common the self-fulfilling cases are. For the Error-Theory Solution to be plausible, our heuristic must plausibly satisfy what we'll call the 'value-requirement': the advantages of employing our heuristic are significant enough to outweigh the costs of doing so. If it's plausible that there's a plentitude of ordinary cases where our heuristic errs (according to CEUT), this fact should cast serious doubt on the claim that the value-requirement is satisfied. So to evaluate whether the value-requirement is satisfied, let's consider some more mundane cases. We'll begin by shifting from the classroom to the workplace:

Tough Day at the Office. Friedrich is anxious about how the day will turn out. He has kept records, and knows that there's about a 40% chance that a given workday will be terrible. Sometimes the workdays are terrible because of how much work he has to do, sometimes due to the commute, and sometimes due to how his boss is

¹³ See Pettigrew (2018, p. 10), Konek & Levinstein (2019, p. 85), and Joyce (2018, p. 252). Although they all describe the case as either "pathological" or "fantastical", they don't precisely specify the reason for describing it so.

treating him. But sometimes, it's just because he's certain that the day will be terrible: whenever he is, this thought makes him depressed, which in turn makes the day terrible. Unfortunately, even if he's certain that the day won't be terrible, there's still a significant chance that it will be terrible.

Friedrich is unfortunate in having a job where 40% of workdays are terrible. But if he becomes certain that he'll have a terrible day, this credence will be maximally accurate. So according to CEUT, epistemic rationality prohibits him from following his evidence in having a moderate credence that the day will be terrible. Moreover, CEUT requires him to be certain he'll have a terrible day!

Not only can our minds affect our quality of life in the workplace, but they can also affect the quality of our athletic performances. In fact, there's a whole field dedicated to optimizing athletes' performance solely using their minds. Consider the following example from basketball:

Missing Mechanics. Markelle just caught a pass and is wide open in the corner with a couple of seconds left on the shot clock. As she elevates to shoot, she knows that she has about a 30 percent chance of making it and that even if she's extremely optimistic, her optimism will barely increase her chances. She also knows that occasionally, before she takes a shot she becomes overrun by pessimism and becomes sure she'll miss. Whenever this happens, she adjusts her mechanics, which results in her missing the shot.

According to CEUT, epistemic rationality prohibits Markelle from following her evidence in only having a modest credence that she'll miss the shot. Furthermore, CEUT demands that she be certain that she'll miss the shot.

Not only can our minds have a powerful effect on our athletic performances, but they can also create strong incentives for us. Consider the following example:

TGIF. Ross and Rachel are thankful once Friday comes around since they eat out or make a nice meal at home. Tonight is Ross' turn to decide where they will go, if he decides to eat out. Since he's usually not in the mood for burgers, he knows it's unlikely they'll go to Bob's Burgers. That said, he hasn't ruled out this option. He knows that if they were to go there, they would have to put their order in far in advance, given the extremely long wait time at Bob's Burgers. So he knows that if he becomes sure they'll go there, he'll immediately put their order in.

Ross' evidence only supports having a low credence in the claim that they'll go to Bob's Burgers. But he knows that if he were to become certain in this claim, he would immediately put their order in, and his credence would turn out maximally accurate. So CEUT prohibits him from following his evidence in being skeptical that they'll go to Bob's Burgers.

Moreover, CEUT demands certainty in this claim.

In addition to our minds creating strong incentives for us, they can also greatly influence the emotional states of ourselves and others.

Reciprocal Empathy. Due to Flor's great love for her daughter Cristina, whenever Flor is certain that Cristina will soon be sad, Flor herself becomes very sad. Cristina feels the same towards her mother, so whenever Cristina sees that her mother is sad, she also becomes very sad. Flor just learned that she has to unexpectedly leave for work in the next few minutes and is about to tell Cristina. Based on the past, Flor knows that there's about a 75 percent chance that Cristina will be sad in the next few minutes upon learning that her mother has to leave.

According to CEUT, Flor would be irrational to not be certain that Cristina will be sad.¹⁴

Clearly, these cases are not bizarre. And they illustrate how easy it is to come up with countless examples where the credences that CEUT recommends diverge from the credences that are intuitively rational. For example, I might have every reason to believe that I'll teach class tomorrow. But if I were certain I'd miss it, then I'd email my students to inform them. And since I'd have no reason to come to class, I'd miss it. I may be unsure whether I'll ever use my new toothbrush again, but if I were sure I wouldn't, I'd throw it away. So, I definitely wouldn't use it again. I may be doubtful that I'll splurge on a spring vacation trip to Aruba, but if I were sure that I would, I would now buy a non-refundable ticket, to get a decent price. Having done so, I would end up splurging. Al the recovering alcoholic might be somewhat uncertain whether he'll relapse by the end of the week, but if he were certain he would, he'd start drinking heavily now. I might be unsure whether I'll procrastinate in finishing a paper for a special edition in a journal. But if I were sure, I'd make sure to at least enjoy my time before rushing to complete it. Having enjoyed this time, I'd have procrastinated. And so on.

¹⁴ Now you might raise the following worry about these cases: do we really find it intuitive that the evidence-disrespecting credences are *epistemically* irrational? Clearly, it's intuitive that they're irrational in some way. But maybe our intuitions are tracking *practical rationality* (or all-things-considered rationality) instead. While some may find it intuitive that the evidence-disrespecting credences are practically irrational, it's also intuitive that they're epistemically irrational. To see this, let's consider a case that disentangles these intuitions. So suppose it's quite unlikely that there will be a really loud noise soon. But it's quite likely that if I were certain that there would be a really loud noise soon, then I would warn my partner by shouting across the room. So it's quite likely that if I were to be certain that there would be a really loud noise soon, then there would be a really loud noise soon. Now is it epistemically rational to be certain that there will be a really loud noise soon? Clearly, it isn't. And given that there are virtually no costs associated with this credence, it doesn't seem to be practically irrational to adopt it. At the very least, the intuition is much weaker that the credence is practically irrational.

The same point holds concerning evidence-respecting credences. Accordingly, it's implausible that our intuitions concerning evidence-respecting credences and evidence-disrespecting credences in the self-fulfilling cases are about practical rationality rather than epistemic rationality.

It's also noteworthy that CEUT has many other counterintuitive verdicts in these self-fulfilling cases. To appreciate this point, let's return to **Painful Prep**. In this case, the intuitively rational credence for Sony to have in the proposition that he'll do badly in the math course is .5. So, it's also intuitively rational for him to have about a .5 credence in the following propositions: he'll not complete the prep work, his GPA will suffer, he'll be upset with his performance in the class, his parents will be frustrated with him, etc. But notice that these propositions all have very high probabilities, conditional on Sony being certain that he'll do badly in the course: he'll not complete the prep work, his GPA will suffer, he'll be upset with his performance in the class, his parents will be frustrated with him, etc. If so, CEUT has the implication that the credal state that is certain in each of them along with the claim that Sony will do badly in Introduction to Proofs has a higher EEU than the intuitively rational credal state that places about a .5 credence in each of these propositions. So there's a *Spillover Problem* for CEUT: not only does CEUT prohibit agents from embracing intuitively rational credences in propositions that they expect to be true if they invest intuitively irrational credences in, but CEUT also prohibits them from embracing many other intuitively rational credences in rationally-related propositions. Furthermore, CEUT requires agents to embrace these intuitively irrational credences in the rationally-related propositions.

Of course, we rarely embrace certainty in these situations. That said, the key point for assessing the Error-Theory Solution is that it's a part of any normal life to regularly be in these situations where your evidence doesn't strongly indicate that some claim is true, but it does strongly indicate that it would be true if you were sure of it. Moreover, the Spillover Problem means that CEUT delivers counterintuitive verdicts about many rationally-related propositions. Given the plentiful number of ordinary cases where CEUT clashes with our

intuitive verdicts, there's a huge cost to employing our heuristic.¹⁵ It seems, then, that the Error-Theory Solution does not satisfy the value-requirement.¹⁶ So while the Error-Theory Solution is a promising solution to the Trade-Off Problem, it struggles with the Problem of Self-Fulfilling Credences.

V. The Epistemic Ratifiability Solution

The final solution we'll consider is James Joyce's *Epistemic Ratifiability Solution*.

According to it, the fact that a credal state (e.g. the intuitively rational credal state in **Imps**) doesn't maximize expected accuracy doesn't mean it is irrational. Epistemic rationality is also a matter of having a usable credal state (Joyce, 2018, p. 253):

The rational epistemic agent who chooses a credal state for her future self will not regard the occupation of that state as the final end, but will see it as an instrument for making estimates, pricing bets, and doing all the other things that credences are supposed to do.

Accordingly, a credal state is irrational only if there's another credal state with a higher expected accuracy that's also *epistemically ratifiable*. Roughly, we can think of a credal state as being epistemically ratifiable for an agent just in case the credal state is most preferable in terms of accuracy, conditional on the agent coming to occupy it. If, conditional on the agent

¹⁵ For a quite different extensional problem concerning causal CEUT, see Caie (2018). Caie argues that in a great number of situations, there's an infinite number of credal states that lack well-defined expected epistemic utilities (EEUs), and that if this is so, causal CEUT delivers incorrect verdicts in these situations. By contrast, I argue that even if every credal state were to have a well-defined EEU in the self-fulfilling cases, both causal and evidential CEUT would deliver incorrect verdicts.

¹⁶ One might object that the evidentialist heuristic actually satisfies the value-requirement. After all, as Pettigrew (2018) notes, the heuristic has always been available, and it's easy to use, too. While this is true, these features don't suffice to make a heuristic a good heuristic—in fact, they count for very little if a heuristic often yields the wrong verdicts. But as we've seen, if CEUT were correct, there would be an abundance of cases in which the evidentialist heuristic would yield the wrong verdicts. For this reason, it seems implausible that a defender of CEUT can plausibly hold that the evidentialist heuristic satisfies the value requirement. Thanks to an anonymous referee for bringing up this concern.

coming to occupy it, another credal state is preferable in terms of accuracy, then the suboptimal credal state is epistemically unratifiable.

We can more precisely characterize the proposal as follows: a credal state C is irrational only if two conditions obtain. First, C has a lower expected accuracy than another credal state C' . Second, the credal state with the higher expected accuracy, C' , is epistemically ratifiable: there is no credal state C'' that disqualifies C' by having a greater expected accuracy, on the condition that C' is adopted.¹⁷ We'll call versions of CEUT that embrace the epistemically ratifiability condition 'R-CEUT'.

But how exactly does R-CEUT help with the Trade-Off Problem? In **Imps**, the intuitively rational credal state is one where Emily is certain that there is now a child before her, and places .5 credence in each of the five propositions stating that a particular summerhouse child will be outdoors in a few minutes. Let us call this epistemically ratifiable state C ; we can represent it as $\langle 1, .5, .5, .5, .5 \rangle$. As we saw, CEUT disqualified C , because it had a lower expected accuracy than a particular intuitively irrational credal state: the one in which Emily is maximally doubtful that there is now a child before her and is certain in the other five propositions. Let us call this second state C' ; we can represent it as $\langle 0, 1, 1, 1, 1, 1 \rangle$. The question then arises: is C' epistemically ratifiable?

It turns out that it isn't. C' does not maximize expected accuracy on the condition that C' is adopted. For on the condition that C' is adopted, there is another credal state that is expectedly more accurate: this is the state which gives maximal credence to all 5 propositions. (Each of the 5 propositions is guaranteed to be true, conditional on Emily adopting C' .) Call this third state C'' ; it can be represented as $\langle 1, 1, 1, 1, 1 \rangle$. C'' shows that C'

¹⁷ While Joyce (2018) isn't entirely clear on which probability function is to be used when calculating the *objective* expected accuracy of C'' on the condition that C' is adopted, it's plausible that the agent's evidential probability function is to be used.

is not epistemically ratifiable. So on Joyce's account, the fact that C' has greater expected accuracy than C does not disqualify C from being rational.

Adding an epistemic ratifiability condition to CEUT results in a promising solution to the Trade-Off Problem. Not only is the epistemic ratifiability condition helpful in this way, it is also a natural understanding of the thought that part of the role of credences is their usability (Joyce, 2018, p. 265). To illuminate the connection between epistemic ratifiability and usability, suppose there's a credal state that maximizes expected accuracy, but isn't epistemically ratifiable. Then, conditional on an agent adopting the given credal state, another credal state is preferable concerning accuracy. And if this is the case, then she can't count on using this credal state to make decisions once she adopts it. If she can't count on using the suboptimal credal state (since she must first shift to an optimal credal state before acting), then it can't disqualify other credal states. If we consider **Imps**, the unusability of $\langle 0, 1, 1, 1, 1, 1 \rangle$ is reflected in the fact that Emily wouldn't be able to rely on her certainty that there's not now a child in front of her when deciding whether to use all of her strength to blindly kick a nearby soccer ball before the summerhouse children come out.

But can the Epistemic Ratifiability Solution also help solve the Problem of Self-Fulfilling Credences? We've seen that Sony being almost certain in (B): he'll do badly in Introduction to Proofs, maximizes expected accuracy. Is this credence in B epistemically ratifiable, though? Notice that it's extremely likely that conditional on him coming to adopt this high credence in B, he won't complete the prep work for the course. And it's also extremely likely that conditional on him not doing the arduous preparatory work, he will do badly in the course. Accordingly, conditional on him coming to occupy the credal state that is virtually certain in B, this credal state is most preferable concerning accuracy. It's plausible, then, that the almost-certain credal state has an unsurpassed expected accuracy, conditional

on Sony coming to occupy it. So it turns out that the almost-certain credal state is epistemically ratifiable.

And if the almost-certain credal state has a higher expected accuracy than the intuitively rational one, and it is also epistemically ratifiable, then R-CEUT prohibits Sony from adopting the intuitively rational one. So even R-CEUT clashes with our intuitive verdicts in self-fulfilling cases.¹⁸ Therefore, although there are some promising solutions to the Trade-Off Problem, the Problem of Self-Fulfilling Credences resists these solutions.¹⁹

VI. The Problem of Diachronic Trade-Offs

We've seen how the future poses a problem for CEUT via our credences about the future. There's still another way in which the future generates a problem for CEUT: the EU of our future credal states. To this end, it will be helpful to consider the following case involving standard decision theory:

Retirement Fund. Evan has amassed a large sum of money with the goal of early retirement. Although he could retire early, he could also rent a lavish yacht for a month and enjoy an incredible vacation. If he were to do so, though, he would not be able to retire early. While he would most prefer to enjoy the incredible vacation if

¹⁸ While Joyce (2019, p. 259) considers a self-fulfilling case from Greaves (2013, p. 916), her case is importantly different from our self-fulfilling cases, as I explain in fn. 8. After all, her case is silent on which credence is supported by the evidence. But in our cases we know that the evidence supports a middling or low credence in the relevant propositions even though the agents know that if they were to have a high credence in the propositions, then these propositions would be true. So even though Joyce acknowledges (and argues) that the self-fulfilling credence in Greaves' case is rational to adopt, the intuitive cost of embracing this claim as applied to our cases is much higher.

¹⁹ For another solution to the Trade-Off Problem, see Joyce & Weatherson (2019). According to their solution, it turns out that certain versions of CEUT are consistent with our intuitive verdict in **Imps**. That said, they concede that different trade-off cases can be constructed where CEUT licenses adopting evidence-disrespecting credal states. Another limitation with their solution is that it only helps with trade-off cases, which the Problem of Self-Fulfilling Credences doesn't rely on.

he knew that he would only live for a month, he most wants to use the money he has saved up to retire early.

Intuitively, it doesn't make sense for Evan to spend all of the money he has saved up on an outlandish vacation. When considering what's rational for Evan to do, we can't merely take into account the value of the present and the near future, we must take into account the value of the whole future too. And when we do, it's clear that it would be irrational for him to spend all of the money he has saved up on a lavish vacation. Evan's case illustrates that if we only consider the present moment and near future, the value that these times have to Evan doesn't settle the question of what action he should perform; we need to know the value of the whole future.

According to standard decision theory, our intuitive judgment concerning Evan's situation is correct. Since standard decision theory takes into account the totality of an agent's preferences, it implies that not spending his retirement money on a lavish vacation has a higher expected utility than spending it. Now if standard decision theory were 'myopic', and didn't take into account the value of the whole future to an agent, it would deliver the incorrect verdict that Evan should take the lavish vacation. This myopic version of decision theory would be implausible precisely because it would be myopic. Accordingly, since practical rationality is teleological and the future can be valuable to an agent, practical rationality's norms should be sensitive to the value of the whole future.

Similarly, if epistemic rationality is teleological and future credal states have EU, then epistemic rationality's norms should be sensitive to future EU. And as Pettigrew says:

AFE [Accuracy-First Epistemology] is a teleological account of the rationality of doxastic items...Roughly speaking, AFE takes a doxastic item to be irrational if it does not promote epistemic goodness. More particularly, it is a teleological account

that is committed to veritism, the thesis that the sole fundamental source of epistemic goodness for a doxastic item is its accuracy.²⁰

So as we've seen from Pettigrew, Konek and Levinstein, and Greaves and Wallace, epistemic rationality *is* teleological and future credal states *do* have EU, according to CEUT.

Nevertheless; CEUT's norms (as they are typically formulated) are only sensitive to proximate EU. According to CEUT, the EEU of credal state C is merely determined by C's accuracy at each possible world, along with the conditional probability of each world. The expected accuracy of the credal states that C will expectedly bring about do not bear on C's EEU. But since future credal states also have EU, how could epistemic rationality's norms fail to be sensitive to the EU of these future credal states? It would be one thing if future credal states didn't have EU, but they do, according to CEUT. So why should epistemic rationality's norms discriminate against future EU? Just as they're not sensitive to *how* or *where* you acquire EU, why should they be sensitive to *when* you acquire EU? Accordingly, it's plausible that well-motivated versions of CEUT are sensitive to future EU.²¹

Accordingly, without getting into the details of how to formally incorporate future EU into CEUT, we can see that if we do it, agents will sometimes be required to accept what we might call a *diachronic* trade-off.²² An agent faces a diachronic trade-off when she can

²⁰ The quotation appears in a previous draft of Pettigrew (2020).

²¹ At the very least, it's plausible that there's a serious explanatory problem here for CEUT. While Joyce & Weatherson (2019) acknowledge and reject versions of CEUT that are sensitive to future EU, they make no attempt to explain how CEUT's teleological foundation is consistent with CEUT discriminating against future EU.

²² While you wouldn't think that diachronic trade-off cases pose a problem for CEUT—since it doesn't take into account future EU—I'm arguing that well-motivated versions of CEUT turn out to be vulnerable to such cases. See: Berker (2013a), (2013b), Firth (1981), and Jenkins (2007) for a similar problem—also using diachronic trade-off cases—for belief versions of epistemic consequentialism.

sacrifice the EU of her initial credal state for the sake of the EU of her future credal states.

Consider the following example:

Too Bad to be True. Tamara is a particle physicist who's aware that there's only one feasible plan for her to discover the answers to a vast range of her research questions. To carry out the plan, she needs to write up an impressive grant proposal for a once-in-a-lifetime research grant that will provide her with the time and resources to answer these questions. As long as she's in a good state of mind and devotes the requisite time, she knows that it's very likely that she'd receive the grant. Unfortunately, she just received some compelling information from the police that her son committed a horrific crime. She knows that if she were to believe this claim about her son, she would fall into despair. And with this despair, she knows she wouldn't be able to write up an impressive grant proposal by the grant's deadline, meaning that she would never answer these research questions.²³

Is it irrational for Tamara to have a high credence in the claim about her son? Given that her evidence strongly supports the claim, the intuitive answer is a resounding 'no.' What does CEUT say? Tamara's case brings out a difference between the verdicts of the myopic and non-myopic versions of it. According to the myopic version, Tamara's having a high credence in the claim about her son is required since the proximate credal state that places a high credence in it has a higher EEU than the credal state that places a low credence in it. The disparity between the EEU of the expected future credal states is irrelevant according to the myopic version.

But matters are quite different for the non-myopic version. According to it, it is irrational for Tamara to have a high credence in the claim about her son, since its norms are

²³ See Nozick (1993, p. 69), for the inspiration for **Too Bad to be True**.

sensitive to her whole epistemic future. Tamara knows that if she were to have a high credence in the claim about her son, she wouldn't be able to write an impressive grant proposal by the deadline. So, she's aware that she'd always have middling credences in the potential answers. If she rejected the claim about her son, though, she could be in a good state of mind, which would allow her to write up an impressive grant proposal. She'd then likely have very accurate future credal states, having high credences in the right answers and low credences in the wrong ones. So it's plausible that CEUT's teleological foundation requires agents to accept diachronic trade-offs. This is the *Diachronic* Trade-Off Problem.

Like the Problem of Self-Fulfilling Credences, the Diachronic Trade-Off Problem seems to pose a more a challenging problem for CEUT and its teleological foundation than the (synchronic) Trade-Off Problem. Tamara's situation is an ordinary one. Many parents have received strong evidence to be quite confident that their children have done troubling things, and know about the emotional costs of believing these claims about their children. Given the ordinariness of Tamara's situation, we can be pretty confident in our epistemic rationality judgments concerning it. As we've mentioned before, this stands in stark contrast to the bizarre nature of **Imps**.

VII. The Error-Theory Solution Reconsidered

Since we're dealing with another extensional problem, it's worth considering whether the two promising solutions to the Trade-Off Problem may be more helpful this time around. So let's begin with the Error-Theory Solution. When we previously considered it, the crucial question was whether the value-requirement was satisfied. This question is pertinent once again. If there is a plentitude of cases where our heuristic clashes with CEUT's verdicts, then it's very unlikely that the value-requirement is satisfied.

As argued above, our heuristic diverges from the myopic version of CEUT in the self-fulfilling cases. Given the abundance of these cases, there are a substantial number of cases where our heuristic implies the incorrect verdict (according to CEUT). But does the non-myopic version concur with the myopic version in these cases? While they might not agree in all such cases, the expected epistemic futures of the self-fulfilling credal states appear to be just as good (if not better) than the expected epistemic futures of the intuitively rational credal states. So the commonness of the self-fulfilling cases provides a good reason by itself for doubting that the value-requirement is satisfied.

The prevalence of diachronic trade-offs provides another one. Let's examine some of these cases.

You Got to Believe. Brian is a pesticide researcher for the government. His research has come to a halt the past few years as he's come to doubt that the government is only using his research for benign purposes. He's aware that his low credence in this proposition is the reason he's been unmotivated to research. He knows that if he were to become certain in this proposition, he would regain his motivation. Further, he's aware that if he were to possess this motivation, he would be able to discover the answers to a multitude of his research questions.

Brian's evidence supports having a low credence in the proposition that the government is only using his research for benign purposes. But if he were to reject his evidence and be sure of this proposition, then he'd have a lot of motivation to carry out his pesticide research project for the government. And he knows that having carried out this project, he could attain much higher credences in the correct answers to his research questions, and much lower credences in the incorrect answers. According to a non-myopic version of CEUT, Brian would be irrational to respect his evidence.

The next case illustrates how our proximate credal states and desires can bear on the accuracy of our future credal states via our future environments.²⁴

The All-Consuming Fad. Michael desperately wants to lose weight. He's tried almost every diet, but hasn't found any success. Each time he's learned of a new fad diet, he's become quite confident in a bunch of nutritional claims that he later realized were dubious. Given his track record, he knows he's quite gullible. He recently heard of the new carnivore diet, an extreme diet that requires you to only eat meat. He's actively trying to avoid learning about the diet, but knows that he could easily get suckered in if the opportunity arose. He recently saw a flyer saying that the famously charismatic inventor of the diet, Dr. Dawn Aker, would be giving a nearby talk tomorrow night. Michael knows that if he were confident that Dr. Aker would be giving a nearby talk, then he would very likely attend.

On the one hand, if Michael respects his evidence concerning the location of Dr. Aker's talk, he can expect to have a poor epistemic future. On the other hand, if he rejects his evidence, he can expect to preserve the quality of his epistemic future by avoiding this dubious fad. According to non-myopic versions of CEUT, Michael would be irrational to respect his evidence concerning Dr. Aker's talk.²⁵

The value-requirement of the Error-Theory Solution looked to be in sufficient trouble to defend myopic versions of CEUT due to the prevalence of the self-fulfilling cases as illustrated by **Painful Prep, Tough Day at the Office, Missing Mechanics, TGIF, Reciprocal Empathy**, along with the others. It looks to be in even worse shape to defend

²⁴ See Berker (2013b, p. 372, fn. 16), for the inspiration for **The All-Consuming Fad**.

²⁵ Encountering people such as Dr. Aker is one way that we can expectedly worsen our epistemic futures. We can also do so through encountering putative gurus concerning politics, religion, self-help, and personal finance.

non-myopic versions of CEUT due to these self-fulfilling cases along with the frequency of diachronic trade-offs as evidenced by **Too Bad to be True, You Got to Believe**, and **The All-Consuming Fad**.

VIII. Reconsidering the Epistemic Ratifiability Solution

While the Error-Theory Solution to our second problem doesn't seem promising, there's still the Epistemic Ratifiability Solution to consider. According to R-CEUT, agents are only required to adopt a particular credal state if it is *epistemically ratifiable*: the credal state is most preferable in terms of accuracy, conditional on the agent coming to occupy it. While diachronic trade-offs pose a problem for non-myopic versions of CEUT, non-myopic versions of R-CEUT do not require agents to accept such trade-offs. Since the intuitively irrational credal states in diachronic trade-off cases are not epistemically ratifiable, these credal states do not disqualify the credal states with lower EEU's.

To see that these credal states are not ratifiable, let's reconsider **Too Bad to be True**. In this case, Tamara's evidence makes it extremely likely that her son has committed a horrific crime. She knows that if she places a low credence in this claim, she can persist in a good state of mind. And being in a good state of mind, she can write up an impressive grant proposal by the grant's deadline, which will allow her to discover the answers to a vast range of her most important research questions. So she knows that by embracing a credal state with a low credence in the claim about her son, she can have a much better epistemic future than if she embraces a credal state with a high credence in it. But according to R-CEUT, this (intuitively irrational) credal state is not epistemically ratifiable. For Tamara, conditional on her coming to occupy a credal state with a low credence in the claim that her son committed a horrific crime, a credal state with a higher credence in this claim is preferable in terms of

accuracy. Since a credal state with a low credence in this claim is not most preferable in terms of accuracy, such a credal state is not epistemically ratifiable.

While the Epistemic Ratifiability Solution to the Problem of Diachronic Trade-Offs is intriguing, it suffers from a serious difficulty: the motivation for R-CEUT makes R-CEUT ad hoc, R-CEUT becomes problematically myopic in another way. To see this, consider the motivation for epistemic ratifiability (Joyce, 2018, pp. 261-262, my emphasis):

As stressed above, when an epistemic agent chooses a credal state she is not choosing a final end. She is choosing the beliefs she will use, starting at t , *to confront the future*, process new information and estimate truth-values and other quantities. *Real credences are essentially forward looking: by choosing them one is choosing to live with their downstream effects.*²⁶

Or as Joyce says elsewhere (2018, p. 253, my emphasis): “The rational epistemic agent...chooses a credal state for her *future self*.” Let’s suppose that credal states have this purpose of being essentially forward looking. If so, a natural question arises: why is epistemic rationality only sensitive to the usability of the credal state one chooses, and not sensitive to the usability of one’s expected future credal states? Why is epistemic rationality myopic concerning the usability of credal states?

Just as CEUT’s teleological foundation implies that standard versions of CEUT are problematically myopic concerning EU, R-CEUT’s foundation reveals that R-CEUT is problematically myopic concerning the usability of credal states. This can be illuminated if we consider the following quote from Joyce (2018, p. 256): “The choice of a credal state at t_0 is merely a means to the end of producing credences that the chooser will be happy to *use* at

²⁶ While Joyce discusses a related notion (being a ‘sham’ credence), the fundamental notion behind epistemic ratifiability is usability.

t.” In Joyce’s terminology, *t* is just the time at which the credal state is adopted.²⁷ But why should epistemic rationality only factor in whether an agent is initially happy to use the credal state? Suppose that epistemic rationality aims at acquiring EU, where its norms are sensitive to both proximate and downstream EU. If so, and epistemic rationality also aims at having usable credal states, shouldn’t its norms be sensitive to both proximate and downstream usability? And if the usability of future credal states is important to epistemic rationality, it looks like Tamara is rationally required to accept the diachronic trade-off. Tamara knows that if she accepts the trade-off, the high and extremely accurate credences of her future credal states will serve her much better in making decisions and for making accurate truth-value estimates than the middling and much less accurate credences. As Joyce says (2018, p. 257):

When assessing a believer’s credences...we want to know how well or poorly the credences will serve her when she uses them as the basis for making accurate truth-value estimates.

If we consider the combined accuracy and usability of her credal states, it’s plausible that accepting this trade-off will result in an overall better epistemic life. So if the motivation to modify CEUT stems from the usability of credal states, a well-motivated version of CEUT will deliver the verdict that it’s sometimes rational to accept diachronic trade-offs.

To avoid this problem for R-CEUT, this motivation stemming from the usability of credal states could be rejected. But doing so would create a new problem; R-CEUT would become unmotivated. So just as the Epistemic Ratifiability Solution and the Error-Theory

²⁷ Joyce uses t_0 to refer to a previous time where the agent can choose the credal state she will have at *t*.

Solution can't help with the Problem of Self-Fulfilling Credences, they're also unhelpful with the Problem of Diachronic Trade-Offs.

IX. Diagnosing the Problematic Future

The arguments in the previous sections provide good reason for thinking that the future poses a couple of significant problems for CEUT. When we're considering what's rational to think about the future, CEUT often delivers the wrong verdicts. It also frequently delivers mistaken verdicts once CEUT's norms become sensitive to future EU. These problems raise a natural question: what's the fundamental source of them? A natural answer: CEUT's (foundational) teleological assumption that epistemic rationality's norms are instrumental norms that serve epistemic rationality's aim.

Now you might think that the aggregational aspect of CEUT is the culprit, rather than CEUT's teleological assumption. According to both the myopic and non-myopic versions of CEUT, the epistemic rationality of a particular credence depends on whether the credence is a part of a credal state that maximizes EEU. Even if a particular credence maximizes EEU when considered in isolation, it may not be part of a credal state that maximizes EEU. This aggregational feature of CEUT is certainly part of the explanation for its incorrect verdicts in diachronic trade-offs such as **Too Bad to be True** and synchronic trade-off cases like **Imps**. In these cases the EEU of a particular credence can be sacrificed to occupy a credal state with an unsurpassed EEU. That said, the aggregational feature of CEUT can't be the primary source of its problems, since it's explanatorily irrelevant in self-fulfilling cases like **Painful Prep**.

These cases do not involve sacrificing the EEU of a particular credence for the sake of a credal state with an unsurpassed EEU. In these cases we can just focus on the EEU of credences in a particular proposition. The agent's evidence supports a middling credence in

that proposition, but also makes it extremely likely that if the agent were to adopt an extreme credence in the proposition, it would be true. Since an evidence-disrespecting credence has a higher EEU than the evidence-respecting one, the evidence-respecting credence is epistemically prohibited.

Another culprit to consider is CEUT's theory of EU. While the standard version of CEUT say that EU is just accuracy, you might instead think that EU is just matching the evidence. But to see that CEUT's theory of EU is not the culprit, let's consider the self-fulfilling cases again. On a natural way of thinking about these cases, you expect that the extreme self-fulfilling credences will perfectly match the evidence as soon as they're adopted. So an evidential-matching-centered theory of EU won't help with the Problem of Self-Fulfilling Credences. Furthermore, CEUT's revised teleological assumption that epistemic rationality aims at matching the evidence would still motivate taking into account future evidence-matching, preserving the Diachronic Trade-Off Problem. So it doesn't look like CEUT's theory of EU is the source of the problems.²⁸

Self-fulfilling cases sharply illustrate the distinction between whether a credence is supported by your evidence and whether your evidence suggests it will be accurate, if you adopt it. While the two properties are typically correlated, they come apart in self-fulfilling cases. In these cases, CEUT prohibits you from adopting the evidence-respecting credence and directs you to adopt the extreme credence that will tremendously raise the probability that it is perfectly accurate. Supposing that epistemic rationality aims at acquiring accuracy, it would make sense that epistemic rationality would take into account how accurate a credal

²⁸ Berker (2013a) and Talbot (2014) also argue that the source of the counterintuitive verdicts of views like CEUT (e.g. belief versions of epistemic consequentialism) is not their theories of EU. While Berker and I think this gives us a good reason to reject such views, Talbot argues that we're committed to accepting these counterintuitive verdicts.

state is, if you adopt it. After all, if epistemic rationality aims at acquiring accuracy, how could its norms fail to be sensitive to how adopting a credal state bears on its accuracy? This is why the fundamental problem with CEUT is its teleological foundation.²⁹

While CEUT represents a natural way to understand the close connection between epistemic rationality and accuracy, the future reveals a couple of problems for this way of understanding the connection. Practical rationality is clearly teleological, and this fits well with taking decision theory as an illuminating account of practical rationality. But if the connection between epistemic rationality and accuracy is not teleological, it might be worthwhile to look beyond decision theory for an explanation of that connection.³⁰

Acknowledgements Many thanks to Zach Barnett, Jamie Dreier, Adam Pautz, Richard Pettigrew, an anonymous referee for *Synthese*, and the audience at Pepperdine’s Philosophy Colloquium Series in 2017. I’m especially grateful to David Christensen and Josh Schechter for their many helpful suggestions.

²⁹ Thus the problems we’ve been examining relate to another sort of problem that’s been discussed in the literature. As others have pointed out, if epistemic rationality is teleological, then epistemic rationality should apply to non-doxastic acts such as eating sandwiches or building the Large Hadron Collider, since such acts can be expected to do better or worse in terms of acquiring EU. This worry concerning CEUT also seems to come from CEUT’s teleological foundation. While the problems may be similar in this way, some might reply to the ‘sandwich’ problem by claiming that it is definitional of epistemic norms that they only apply to doxastic options, not to lunch options. This reply is of no help in solving the problems in this paper, since the problems concern doxastic options. See Greaves (2013, p. 922), for this “definitional” response. For more on this worry concerning teleological theories of epistemic rationality, see: Arpaly (unpublished Ms.), Horowitz (2019), and also Konek & Levinstein (2019, p. 100).

³⁰ See Marxen (2021) for an alternative account of the connection and for a different argument—which appeals to a connection between epistemic rationality and practical rationality—for the conclusion that potentially self-fulfilling credences are not thereby rational to have.

References

- Anscombe, G.E.M. (1957). *Intention*. Oxford University Press.
- Berker, S. (2013a). Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122(3), 337-393.
- Berker, S. (2013b). The rejection of epistemic consequentialism. *Philosophical Issues*, 23(1), 363-387.
- Caie, M. (2013). Rational probabilistic incoherence. *Philosophical Review*, 122(4), 527-575.
- Caie, M. (2018). A problem for credal consequentialism. In J. Dunn & K. Ahlstrom-Vij (Eds.), *Epistemic consequentialism* (pp. 182-219). Oxford: Oxford University Press.
- Carr, J. (2017). Epistemic utility theory and the aim of belief. *Philosophy and Phenomenological Research*, 95(3), 511-534.
- Firth, R. (1981). Epistemic merit, intrinsic and instrumental. *Proceedings and Addresses of the American Philosophical Association*, 55, 5-23.
- 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.
- Horowitz, S. (2019). Predictably misleading evidence. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-order evidence: New essays* (pp. 105-123). Oxford University Press.
- Jenkins, C. (2007). Entitlement and rationality. *Synthese*, 157(1), 25-45.
- Joyce, J.M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575-603.
- 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.
- Joyce, J.M., & Weatherson, B. (2019). Accuracy and the imps. *Logos and Episteme*, 10(3), 262-282.
- Konek, J., & Levinstein, B. (2019). The foundations of epistemic decision theory. *Mind*, 128(509), 69-107.
- Nozick, R. (1993). *The nature of rationality*. Princeton: Princeton University Press.
- Marxen, C. (2021). On the epistemic rationality and significance of self-fulfilling beliefs. *Synthese*.
- 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.
- Pettigrew, R. (2020). Logical ignorance and logical learning. *Synthese*.

Talbot, B. (2014). Truth promoting non-evidential reasons for belief. *Philosophical Studies*, 168, 599-618.