



A Dilemma for Higher-Level Suspension

ABSTRACT: *Is it ever rational to suspend judgment about whether a particular doxastic attitude of ours is rational? An agent who suspends about whether her attitude is rational has serious doubts that it is. These doubts place a special burden on the agent, namely, to justify maintaining her chosen attitude over others. A dilemma arises. Providing justification for maintaining the chosen attitude would commit the agent to considering the attitude rational—contrary to her suspension on the matter. Alternatively, in the absence of such justification, the attitude would be arbitrary by the agent’s own lights, and therefore irrational from the agent’s own perspective. So, suspending about whether an attitude of ours is rational does not cohere with considering it rationally preferable to other attitudes, and leads to a more familiar form of epistemic akrasia otherwise.*

KEYWORDS: epistemic akrasia, higher-order evidence, suspension of judgment, normative uncertainty

1. Higher-Level Suspension

Some combinations of doxastic attitudes seem to entail irrationality. If we are more confident that Linda is a feminist bank teller than we are that she is a bank teller, then at least one of our attitudes is irrationally had. The same can be said of agents who strongly believe that bad things happen only to those who have done wrong, strongly believe that a bad thing has happened to them, but strongly disbelieve that they have done wrong. The thought is that no situation rationally permits such combinations of attitudes, and so an agent must be committing some error of rationality by having them.

Among the candidates for irrational combinations are what have been called *epistemically akratic* attitudes. These are combinations of certain *lower-level* and *higher-level* doxastic attitudes—where a lower-level attitude is about some proposition *p*, and a higher-level attitude is about that lower-level attitude. Epistemically akratic combinations involve a lower-level doxastic attitude toward *p* and a higher-level attitude of doxastic disapproval (like disbelief or low confidence) toward the rationality of that lower-level attitude. For example, an agent has epistemically akratic attitudes if she believes both that civilization hardly contributes to global warming and that this belief of hers is irrational. Such an agent would not be endorsing her own view about global warming as rational,

Special thanks for feedback on this project go to Sara Aronowitz, Lisa Benossi, Sven Bernecker, Sofia Bokros, Jennifer Rose Carr, Anna-Maria A. Eder, Thomas Grundman, Giulia Napolitano, Francesco Praolini, Luis Rosa, Mattias Skipper, Paul Silva, and Elise Woodward.

which would make it unclear why she continues to have it. So, it is natural to think that something is rationally defective with agents who have epistemically akratic attitudes. Of course, not all combinations of lower-level and higher-level doxastic attitudes suggest irrationality. *Epistemically enkratic* combinations of attitudes, which are had when agents consider their own beliefs to be rational, are unproblematic.

Between clearly akratic and clearly enkratic combinations of doxastic attitudes there is interesting territory. There, we find agents who neither believe nor disbelieve that their own beliefs are rational. It is harder to know what to say about such agents and their beliefs. As Roderick Chisholm (1989: 6) states, while it is wrong to reason from presuppositions that we explicitly consider unjustified, there is nothing clearly wrong with reasoning from presuppositions that we merely do not explicitly consider justified. So there is at least a *prima facie* difference between the rational status of combining a belief with disbelief that it is rational, and the rational status of combining a belief with suspension of judgment about whether it is rational. In what follows, I describe agents who suspend judgment about whether a particular belief of theirs is rational as being in the state of *higher-level suspension*. My focus is the rational status of such a combination of attitudes—that is, having a belief and at the same time suspending judgment about whether that belief is rational. In question form:

Higher-level suspension question. Is it ever rational to suspend judgment about whether a particular belief of ours is rational?

Quite a bit rests on our answer to the higher-level suspension question. On the one hand, if higher-level suspension is often rational, peer disagreement might not have as strong an effect as some think it has (Feldman 2006; Christensen 2007; Cohen 2013). We could believe *p*, meet an epistemic equal who disbelieves *p*, and continue believing *p* while suspending about whether our belief is rational. On the other hand, suppose that upon consideration of any proposition *p*, we should believe, disbelieve, or suspend judgment about it (Smithies 2012: 275; Turri 2012: 356). And suppose that akratic combinations are irrational (contra Weatherson 2019; Lasonen-Aarnio 2020). If higher-level suspension is irrational, too, then whenever we should believe *p*, we should also believe that we should believe *p*. Our justification for any proposition would thus turn out to be self-intimating, thereby establishing a surprisingly strong level-connection principle (for discussions of level-connection principles and akrasia, see Smithies 2012; Dutant and Littlejohn 2016; Kiesewetter 2016). So interesting things follow both from a positive and from a negative answer to the question.

Here I pursue a (qualified) negative answer. I argue that it is irrational to suspend judgment about whether a particular belief of ours is rational, on the assumption that outright akratic combinations of attitudes are irrational. My goal is to show that if epistemic akrasia is irrational, so is higher-level suspension. Suspending about whether our belief is rational is no better than believing that it is not. If this is right, it would mean that denying the rationality of akrasia has more radical implications than it seems.

Two preliminary clarifications. First, the sense of *suspension* I appeal to is that of *medium confidence*, or a *middling credence* (Hájek 1998: 204; Christensen 2009:

757). On this take, suspending about whether a proposition is true amounts to being roughly equally confident that it is true and that it is false. Although far from being the only option, it is natural to understand *suspension* in this way. Paradigm cases of suspension are ones where we think that the proposition in question is about as likely as its negation. We suspend about whether a coin will land heads when flipped. We suspend about whether each of two evenly matched teams will win the game. So this medium confidence sense of *suspension* looks like a decent place to start when considering the rational status of higher-level suspension. But more importantly, other takes on *suspension* would not work well in this context. Consider, for instance, understanding the notion as *openness to inquiry* (Friedman 2017), or as a kind of *doxastic abstinence* (Crawford 2004; Boghossian 2008). Granted, it is typically rationally permitted to be open to inquiry, or to have no view about matters that we have not considered (though if morality encroaches on rational permission, that might render some inquiries or avoidance of consideration irrational; see Gardiner 2018; Moss 2018; Basu 2019). But these takes on *suspension* leave a key issue in a number of related debates about epistemic akrasia untouched. The issue is that of how confident or uncertain we may rationally be about the rationality of our attitudes. For while it may be rational to be open to inquiry about whether our attitudes are rational, or to have no view about whether they are rational, the question of how confident such inquiring or abstaining agents should be about the rationality of their attitudes would remain. And it is the answer to this question that would shed light on matters of high epistemological significance, like the normative impact of higher-order evidence such as peer disagreement, and relatedly, whether we have any a priori justification regarding what rational requirements obtain. These are matters that we could not tackle by reading *suspension* in alternative ways. So the arguments to come will also not treat *suspension* as compatible with high/low credence (Buchak 2014; Staffel 2015; Jackson 2020), or as the reflective belief that our evidence is inconclusive (Rosenkranz 2007; Raleigh 2021), but should apply to some views of suspension as an imprecise credence (Joyce 2010; Sturgeon 2010)—as long as the relevant credence interval is around a middling credence.

Second, the term *rational* that features in the higher-level suspension question can be read *propositionally* or *doxastically* in each of its two appearances. This means that there are in principle four higher-level suspension questions to distinguish here. I focus on the doxastic/propositional reading—that is, where the first instance of *rational* is read doxastically and the second propositionally, as other discussions of epistemic akrasia seem to (see Christensen 2013; Horowitz 2014; Schoenfeld 2015: they discuss whether we may rationally disbelieve [doxastic justification] that our beliefs are supported by the evidence [propositional justification]). The resulting and more precise higher-level suspension question that follows from these two clarifications is this:

Higher-level suspension question.* Can we ever rationally have both a doxastic attitude *a* toward a proposition *p* and a middling degree of confidence that our evidence rationally justifies *a*?

As with the question of whether epistemic akrasia can be rational and with arguments that it cannot, below I focus on potential conflicts within the agent who is in the state of higher-level suspension (see Greco 2014; Horowitz 2014). Thus, my arguments directly concern structural rationality, or coherence, as opposed to substantive rationality, or what the evidence supports (Worsnip 2018 distinguishes *irrationality* from *incoherence*). With that, I now explore the negative answer that I am after.

2. The Negative Answer So Far

Admittedly, the view that higher-level suspension is never rational faces some intuitive pushback. It can be hard to tell what rationality permits in different situations, and whether our attitude falls within the permitted range. This seems obviously true when the available evidence is complex, or when there is considerable disagreement over what rationality permits. In general, suspension of judgment is called for when we should not consider a proposition much more likely than its negation. It is therefore tempting to think that as long as it is hard to tell whether our attitude is permitted, we may suspend about whether it is. Such suspension can appear rational even when our attitude does fall within the rationally permitted range. But if this is right then the negative answer is wrong. Moreover, since what rationality permits *is* a complex and highly contested matter, it is frequently hard to tell whether our attitudes are rational. So we might expect higher-level suspension not to be just occasionally permitted, but to be frequently. If this is right then the negative answer is dead wrong.

Nevertheless, there are reasons to think that higher-level suspension cannot be rational. In particular, I find two of the considerations in favor of this view especially instructive. The first comes from Michael Bergmann, who says that suspension about whether one's attitude is rational defeats one's justification for that attitude: 'If you are considering whether the actual basis of your belief that *p* is indicative of *p*'s truth and you find yourself resisting the belief that it is (because you have considered the matter and you have no idea whether it supports *p* or not), that seems to undercut your justification for believing *p* in the same way as if you believed outright that the actual basis for your belief that *p* did not indicate *p*'s truth' (Bergmann 2005: 426). The thought seems to be that having sufficient doubt about whether we believe rationally is ultimately as problematic as outright epistemic akrasia. If high confidence that we irrationally believe *p* undercuts our justification for *p*, then perhaps having no idea whether we believe *p* rationally could do so as well.

The second consideration comes from Michael Huemer (2011) and Declan Smithies (2012), who offer similar arguments against higher-level suspension. Huemer takes suspension about a proposition *p* to license assertions like *it may or may not be the case that p*. Smithies takes suspension about *p* to license assertions like *it is an open question whether p*. Accordingly, if we suspend about whether our belief that *p* is rationally justified, we could assert that *we may or may not be justified in believing p*, and that *it is an open question whether we are justified in believing p*. Both kinds of assertions can appear to conflict with asserting *p* at the

same time, as they would give rise to Moore-paradoxical-sounding utterances like *p*, but *we may or may not be justified in believing p*.

Allan Hazlett (2012) gives voice to the opposition. In response to Bergmann, Hazlett asks for more than an appeal-to-intuition that suspension at the higher level defeats our justification for the lower-level attitude in question. In response to Huemer, Hazlett challenges both the thought that it is improper to assert claims like *p*, but *we may or may not be justified in believing p*, and the thought that inability to assert properly a conjunction implies inability to believe it rationally.

I share the impression that the arguments for the irrationality of higher-level suspension do not go far enough. They do, however, establish a key point. They show that something quite confusing is going on with agents who suspend about whether a belief of theirs is rational. So even if the initial confusion regarding agents in this state does not imply their irrationality, we need some story by way of explaining how their chosen view fits with their suspension. Specifically, we need some story by way of explaining why their chosen view is somehow rationally preferable to others. Without such a story, I argue, that chosen view would be arbitrary by these agents' own lights. So, to advance the rational impermissibility of higher-level suspension, we need to show that there is something wrong with different ways of telling that story.

The plan for what follows is to demonstrate both the irrationality of agents who do not have such a story and the irrationality of agents who do. Here is the basic idea. An agent who suspends about whether her attitude is rational must be able to address the concern that she chose the attitude arbitrarily, and thus irrationally. Not addressing the concern would leave the attitude looking arbitrary by the agent's own lights. In that case, the agent would be as irrational as an akratic agent—vindicating Bergmann's suspicion. But the more promising alternative of addressing the concern would not end well either. For doing so would require the agent to defend the rational permissibility of her chosen attitude. In that case the agent would be incoherent, since her suspension about whether her chosen attitude is rational would conflict with her justifying the choice as rationally permissible.

3. Against Higher-Level Suspension: A Dilemma

An agent who suspends about whether the evidence justifies her attitude toward *p* needs special reason to maintain that attitude over others. Without such reason, and given her suspension about the attitude's rationality, the agent would be committed to thinking that she chose the attitude arbitrarily, and that she has it irrationally. So being sufficiently skeptical that an attitude of ours is justified carries with it a unique burden. In a slogan: *with great doubts that an attitude is justified comes great responsibility to justify having it*.

Imagine how a conversation with such an agent might go, depending on whether she had special reason for maintaining her chosen attitude over others. Upon learning that we face an agent who believes *p* but is quite skeptical that her evidence justifies the attitude, we may naturally ask what reasons she has to

maintain her attitude toward p rather than some other attitude. The agent could respond in one of two basic ways, *good answer* or *bad answer*:

1. Good answer. *I have my chosen attitude for good reasons r_1, r_2, \dots, r_n .*
2. Bad answer. *I lack good reason to maintain my chosen attitude.*

Different versions of good answer could help us understand why the agent has her attitude and not another. I discuss those versions soon. However, no version of bad answer would do. Bad answer would make the agent seem patently irrational: she would have an attitude, and maintain it despite conceding that she has it for no good reason. In other words, the agent would be committed to thinking that her chosen attitude is arbitrarily had.

Notice that a bad answer does not imply that the agent is explicitly akratic. The agent who suspends at the higher level and gives a bad answer does not believe that the evidence does not rationally justify her attitude. Instead, the agent considers her attitude to be arbitrarily had, regardless of what the evidence justifies. But if akrasia is an irrational state to be in, so is this. An agent who maintains an attitude despite thinking that she has it for no good reason seems just as irrational as one who maintains the attitude despite thinking there is no good reason for it. Both agents would reason with and act on the basis of their lower-level attitude, and at the same time deny that what they are doing is rational, or advise others not to follow suit. The same type of incoherence or fragmentation that is present in explicitly akratic agents is also present in agents who hold attitudes that they consider arbitrary.

The impression of irrationality that we get here is more robust than the one Huemer (2011) and Smithies (2012) invite us to have. In their discussions, we are asked to view as irrational an agent who asserts p and *my evidence may or may not support p* . But one thing that stands in the way of viewing such an agent as irrational is the possibility that the agent has some special reason for having her chosen attitude. Bad answer eliminates that possibility, and paves the way for concluding that the agent indeed has an irrational combination of attitudes. So, the state of higher-level suspension is irrational when paired with a bad answer of some kind.

It is worth noting that it is not necessary for the agent to be able to express her reasons for maintaining her belief. What is necessary is for the agent to *possess* reasons that she could in principle appeal to by way of justifying her chosen attitude, despite her doubts that the evidence justifies that attitude. In the absence of such reasons, the agent would be committed to considering her belief arbitrary. So, to see if the agent could rationally maintain an attitude while suspending about its rationality, we should look into more promising responses along the lines of good answer.

What could the requisite good answer look like? Perhaps the most tempting answer is that the agent's evidence itself may act as the reason that she has her chosen attitude rather than some other one. But that will not work. Recall that the agent suspends about whether the evidence justifies her attitude. The agent cannot

appeal to the evidence as the reason why she has her chosen attitude, and at the same time be on the fence about whether the evidence justifies that attitude. That would be incoherent.

Consider, by analogy, an inept baker who cannot recall which of two steps of a cupcake recipe comes first: the baking of the mixture, or the application of the frosting. As a result, the baker suspends about which comes first. Suppose that the baker nevertheless proceeds to bake the mixture first. The baker could not justify the action by appeal to the recipe itself, for they are clueless as to which step comes first. So even though the recipe in fact requires baking the mixture first, that fact is not available to the baker as a reason for the action, since they are on the fence about which step comes first. Similarly, an agent who suspends about whether the evidence rationally justifies her attitude needs to be able to appeal to something other than the evidence itself in order to maintain her attitudes rationally.

Perhaps a better good answer could appeal to the attitude that the agent takes her situation *probably* to justify. There is nothing obviously incoherent with being uncertain about which of a few possible attitudes is rationally permitted, and at the same time thinking that one of them is more likely to be permitted than the others. For example, we may think that there is a 24 percent chance that credence c_1 is permitted, a 25 percent chance that credence c_2 is, and a 51 percent chance that credence c_3 is. In such a case, the agent could say that she has c_3 because it is most likely the permitted credence, even though she suspends on whether c_3 is permitted. Call this the *probability answer*.

The probability answer does a good job of accounting for the agent's having the attitude that she has and not another. It identifies a feature of the attitude, namely, the agent's belief that it is more likely than others to be rational, and offers that as the reason to have the attitude. The answer also appears to fit nicely with the agent's middling credence that her attitude is rational. However, the answer ultimately does not fit with suspension about whether the attitude is rational. The problem is that an agent who offers such an answer is implicitly committed to a principle that tells us what we may believe in situations like hers. The probability answer would commit the agent to something like the following principle:

The probably rational principle. When attitude a is probably the rationally permitted one, then it is rationally permitted to have it.

The implicit commitment comes from the agent's appeal to the claim that her attitude is probably rational as the reason why she chose it. For suppose that the agent offered the probability answer and also rejected the probably rational principle. In that case, we would wonder why the agent thinks that the fact that her chosen attitude is probably rational justifies having it. Thus, to address our question of what reason the agent has for maintaining her chosen attitude properly, a good answer like the probability answer requires (at least implicit) belief in a corresponding principle like the probably rational principle.

Now comes trouble. An agent who subscribes to the probably rational principle thinks that when an attitude is most likely rationally permitted, then *it is* rationally permitted. Such an agent cannot take her attitude to be most likely permitted, and at

the same time suspend about whether the attitude is permitted, on pain of incoherence. Thinking that our attitude is probably rational, and thinking that it is rational to have the attitude that is probably rational, together rule out seriously doubting that our attitude is rational.

We get a similar result on a credence framework. If the agent has a high credence in the claim that her chosen attitude is probably rational, and also a high credence that the probably rational principle is correct, the agent cannot coherently have a middling credence that her chosen attitude is rational. However, if the agent has a not-so-high credence in each (say, 0.7 that her attitude is probably rational and 0.7 in the probably rational principle), then the agent could not use their conjunction by way of a good answer. For the agent would only have a middling credence that the principle applies to her situation so as to make her chosen attitude rationally preferable to others.

The probability answer is one example of a good answer. Yet the conclusion we should draw is more general. The same argumentative moves are available against any other good answer that the agent may offer by way of explaining why she has the attitude that she has. Any such answer would effectively be offering some reason for why it is rational to have the attitude that the agent does, despite her suspension about whether the attitude is rational. This would commit the agent to a corresponding principle that says that *when certain criteria are met, a certain attitude is rationally permitted*. But subscribing to such a principle, while thinking that the relevant criteria are met, implies that the attitude is indeed permitted. So the agent could not suspend about whether her attitude is permitted if she has some kind of good answer to why she has that attitude.

Applying the argument to a few other possible good answers can help demonstrate its generality. For starters, consider a suggestion by David Christensen concerning when epistemic akrasia is rational: 'Epistemic akrasia is not, per se, a problem at all. Thinking that a belief of yours is irrational in a particular way should disturb you—that is, give you a reason to change that belief—only insofar as the particular irrationality indicates that a different belief would have greater expected accuracy' (Christensen 2016: 416). If we apply this suggestion to the case of higher-level suspension, we get an agent who suspends about whether her credence c is rational, and offers Christensen's explanation as to why she nevertheless has c . In so doing, the agent would be committing herself to a principle like the following:

Greatest expected accuracy principle. Unless a different credence than one's own has greater expected accuracy, it is rational to maintain one's own credence.

Recall that the question of interest here is whether an agent who suspends about whether her attitude is rational can coherently appeal to the relevant principle by way of justifying her credence given her higher-level suspension. The question is not whether the principle that the agent appeals to is in fact true. The answer to the former question appears to be negative. If the agent believed the principle to be true, she could no longer suspend about whether her credence is rational. She would be committed to believing that it is rational, since she believes that her

credence satisfies the accuracy criterion, and since she takes the principle to be correct.

Another test case for the argument's generality comes from Richard Feldman's (2006; 2007) discussion of disagreement. As Hazlett (2012: 205) notes, Feldman's argument for the need to suspend judgment when disagreeing with a peer seems to rest on a principle about higher-level suspension. According to that principle, when we should suspend about which attitude regarding p is rational, we should also suspend about p . Now suppose that an agent who suspends about whether her lower-level attitude is rational appeals to such a principle in order to justify her lower-level attitude. At that point the agent could look at her lower-level attitude of suspension about p , see that it is in line with Feldman's view, and come to believe that the attitude is rational. But this would not fit with the agent's suspension about which attitude toward p is rational. Subscribing to such a principle would not allow the agent to continue suspending about what attitude she should have toward p .

For a last test case, consider an agent who justifies her chosen attitude by appeal to a kind of permitted arbitrariness. On such a view, if we rationally think that two attitudes toward p are equally likely to be rationally permitted, then it is rationally permitted to choose one arbitrarily. This would not avoid the problem either. An agent who appeals to this view would be committed to the claim that her arbitrarily chosen attitude is in fact rationally permitted, which would conflict with her suspension about whether her attitude is permitted.

The lesson is that the details of any good answer do not matter much. An agent cannot both engage in justifying her attitude over others, and also suspend about whether the attitude is justified. And since not engaging in this kind of justification (or providing a bad answer) leaves the agent's chosen attitude looking arbitrary by her own lights, higher-level suspension is in a bind. It therefore seems irrational to have an attitude toward a proposition p while suspending about whether that attitude is rational.

4. Justificatory Burdens and Higher-Level Certainty

Two concerns about the argument are worth pausing for. One challenges the claim that an agent in the state of higher-level suspension needs to justify having her lower-level attitude over others in the first place. Another states that the argument implies the irrationality of any degree of uncertainty about whether we believe rationally.

4.1. The Need for Justification

The first concern starts with the observation that we often believe things rationally despite having no story as to why we believe as we do. Perceptual beliefs provide one kind of example. We can justifiably believe that *the person over there looks familiar* without being able to justify why we think that they look familiar. As the famous chicken-sexer example shows (Armstrong 1963: 431), an agent can justifiably believe that *this chicken is male* without being able to specify why she

does not believe that it is female instead. Beliefs held by children provide another example, for they may be unable to provide reasons for their beliefs despite having those beliefs rationally. If this is right, then the mere fact that an agent cannot provide justification for having her chosen view over others need not imply that the view is arbitrary or irrational by the agent's own lights. So, the objection goes, an agent who has no answer for why she maintains her attitude is not automatically committed to anything that is in tension with her suspension about whether her attitude is rational.

Let us grant for the sake of discussion that agents typically do not need to be able to justify having their beliefs over others in order to have those beliefs rationally. Demanding such justification may be asking too much of believers, who on many epistemically unproblematic occasions could not offer it. But the thoughts that we should not ask too much of believers, and that we could consider their beliefs rational even if they cannot justify having them, are much less compelling when we learn that the agents themselves have serious doubts that they believe rationally. A chicken-sexer who doubts that her belief about the bird's sex fits her overall perceptual experience should perhaps double check rather than carry on believing that the bird is male (Schechter 2013 argues for double checking as the responsible thing to do in cases of higher-level uncertainty). The fact that an agent has a skeptical attitude toward her own rationality seems to open the door for us to demand (and for the agent to need) something extra.

But regardless of whether the agent can retain her lower-level attitude rationally without answering this demand, the argument against higher-level suspension would still go through. The argument's conclusion is not that the higher-level suspender who cannot provide a good reason to favor her chosen attitude has *that attitude* irrationally. Rather, the conclusion is that the agent's *combination of attitudes* is irrational—much like in the case of epistemic akrasia. The argument points to a conflict that arises from the agent's holding an attitude and suspending about whether that attitude is rational. In the absence of some kind of good answer, the suspension at the higher level commits the agent to thinking that her lower-level attitude is no better than a guess (even when it is in fact much better). The result is that the agent could not rationally maintain her combination of attitudes, irrespective of whether her lower-level attitude is rationally had. So, the response to the concern here is that while agents who higher-level suspend might retain their lower-level attitude rationally despite having no story as to why they have *it* specifically, their entire doxastic state would be irrational.

Notice that we could not say the same of the chicken-sexer who finds herself without a story as to why she believes as she does. If we could, then once we confronted the chicken-sexer about her inability to justify her belief, the argument would yield the unintuitive result that the chicken-sexer's state on the whole is irrational. Yet there is distance between noticing that we cannot offer reasons for a belief, and thinking that there is a good chance that the belief is irrational. A chicken-sexer might appeal to her perceptual experience, and say that *something about it* justifies her judgment that the bird is male—though she is not sure what. This position could be perfectly rational, and prevent the agent from landing in a state of higher-level suspension. In fact, we do this kind of thing all the time. For

example, we might tell our partner of an especially odd colleague whom they have to meet. The fact that we have a hard time explaining what about a person's behavior makes us think that they are odd typically does not make us think that our belief that they are odd is irrational.

Relatedly, some may think that the choice between good answer and bad answer rests on a false dichotomy. An in-between option is for the agent to either suspend judgment or have no view about whether she has good reason to favor her lower-level attitude over others. That way the agent might appear to avoid both horns of the dilemma. The agent would not be providing some good answer that would commit her to viewing her chosen attitude as ultimately rational. The agent would also not be providing some bad answer that would commit her to viewing her chosen attitude as arbitrary and no better than a guess.

However, it does not make a difference whether the agent provides a bad answer, avoids providing an answer, or suspends judgment about whether her attitude is arbitrary. By having serious doubts that her lower-level attitude is rational, and in the absence of some good answer for having it over others, the agent's attitude already looks like a guess. It is at that point that epistemic damage is done, and a burden is placed on the agent to show otherwise. No further suspension of judgment or abstaining from having a view about whether there is a good reason to prefer that attitude would undo this damage. So, whether the agent explicitly offers a bad answer or offers no answer at all, her attitude would remain arbitrary by her own lights.

Importantly, what I say here should apply equally well regardless of what the agent's lower-level attitude is. For instance, when one thinks that there is about a 50 percent chance that, say, a 0.82 credence in p is irrational, it would seem arbitrary to stick to 0.82. There is clearly nothing unique to 0.82 that makes it not look arbitrary when we are highly suspicious of its rationality. But notice that the same can be said of any other lower-level attitude or credence, including suspension of judgment or 0.5. Granted, suspension of judgment and 0.5 do have a certain unique feature, namely being exactly in between full belief and full disbelief. But this feature does not make having such an attitude any less arbitrary when we think there is about a 50 percent chance that it is irrational. For suppose that the agent who suspends about whether her lower-level suspension (or 0.5) is rational is equally confident that full belief (or some particular high credence) is rational. Such an agent would blatantly be picking one of only two viable options for no reason whatsoever. This shows that the unique features of suspension of judgment do not shield it from looking arbitrary when combined with higher-level suspension. And while it may be that in some cases sticking to suspension at the lower level is not arbitrary for some reason, the agent would have to appeal to such a reason in order to avoid any charge of arbitrariness. In other words, the agent would have to use that reason as part of a good answer to why they maintain their lower-level suspension—at which point the problems of providing such an answer would come in.

Lastly, even if the arguments in this paper were to apply only to agents whose lower-level attitude is not suspension, the implications would still go quite far. We would be left with the result that whenever we should have a somewhat decisive view about p —be it belief, disbelief, or some credence other than 0.5—then on the

assumption that epistemic akrasia is irrational higher-level suspension would be irrational too. So we would end up with an extensive prohibition on suspension about whether we believe rationally.

4.2. Normative Certainty

A second concern is that if the argument that I advance works, then it could also establish the irrationality of the slightest doubt that a belief of ours is rational. The thought is that if serious doubts about our rationality result in an incoherent state, then so might slight doubts. Since we frequently make mistakes in reasoning, having some doubt that we believe rationally on a specific occasion seems perfectly rational. So it would be a dubious upshot of the argument on offer if it deemed irrational all agents who have some small degree of uncertainty that a belief of theirs is rational.

To motivate the worry, here is how we may think that a version of the argument could apply to an agent who has only a slight doubt that she believes rationally. Suppose we face an agent who believes, but is uncertain, that her doxastic attitude toward p is rational. As before, we may wonder why the agent has the attitude despite her uncertainty, and think that the agent must justify having her lower-level attitude given her slight higher-level doubt. In response, the agent might say that it is permissible to have an attitude despite a little uncertainty. At that point the agent would seem committed to the claim that her attitude is rational, and should no longer be uncertain about whether it is. Alternatively, if the agent cannot justify having her chosen attitude over others, we may say that her attitude is arbitrary and irrational by her own lights.

Does the argument commit us to the view that rationality requires certainty about whether we believe rationally? One thing to note is that this worry is not unique to the view that higher-level suspension is irrational. The worry applies more generally to views that consider epistemic akrasia irrational. The thought is that if strong distrust that we believe rationally (of the sort p , but *believing p is irrational*) is irrational, then a little distrust amounts to at least a little irrationality. In credence terms, if having a very high credence that our attitude a is irrational makes maintaining a blatantly irrational, then having a low but non-zero credence that a is irrational makes maintaining a somewhat irrational (Lasonen-Aarnio 2020 notes this worry; for a discussion of solutions, see Dorst 2020; Skipper 2021).

But more to the point, an agent's higher-level belief that her lower-level attitude is rational can fit just fine with *some* uncertainty about whether that attitude is rational. This is a key difference between the agent who is only slightly uncertain that her attitude is rational and the agent who is quite skeptical and suspends judgment about whether it is. The agent who suspends judgment about whether her attitude is rational cannot at the same time have a confident view about what attitudes we should have in cases of such higher-level doubt. But the agent who is only slightly uncertain that her attitude is rational can consistently maintain the attitude, and at the same time commit herself to a principle that says that her attitude is rational. As long as the agent is not certain of that principle, her commitment to the claim that the attitude is rational does not entail certainty that it is.

For example, suppose that *S* believes that global warming is largely manmade, and also believes that her belief on the matter is rational. Nevertheless, *S* has some doubts that her belief is indeed rational—say she has a 0.9 credence that it is, and a 0.1 credence that it is not. Suppose we then ask *S* why she believes that global warming is largely manmade, given that she also has some doubt that this is the rational belief to have. *S* has a perfectly good way to answer our question. *S* could say that she believes global warming is largely manmade because this is what her evidence probably justifies, and she also strongly believes (though not with certainty) that it is rational to believe whatever your evidence probably justifies.

5. Conclusion

It can be rational to believe *p* despite having no view about whether we believe *p* rationally. It can also be rational to believe *p* despite our continued inquiry and deliberation about whether *p*. Problems begin when we believe *p* and have serious doubts that we believe *p* rationally. Having such doubts requires us to do something to restore rational harmony. But once we do that something—for instance, by appeal to principles that imply we may believe *p* despite our doubts—we undermine those doubts.

What does this show? The thought that epistemic akrasia is irrational has more extensive implications than it seems. Michael Titelbaum (2015) has argued that if epistemic akrasia is irrational then false beliefs about what rationality permits are themselves irrational. This is his fixed point thesis: ‘mistakes about the requirements of rationality are mistakes of rationality’ (2015: 253). Titelbaum tries to push the implications further by arguing that the best explanation of the irrationality of akrasia is that we have indefeasible a priori justification for true beliefs about what rationality requires (2015: 276). My arguments here can strengthen the case for a link between the irrationality of akrasia and our having such a priori justification. If the irrationality of akrasia implies the irrationality of higher-level suspension, the only option remaining for an agent who should believe *p* is to believe also that she should believe *p*. We would thus get a surprisingly strong level-connection principle, according to which rational justification for any lower-level attitude can never be had without justification that that attitude is rationally justified. So it would not merely be, as Titelbaum’s fixed point thesis suggests, that we always lack support for false beliefs about what rationality requires. For that could happen if we simply have little evidence about what is rational, and if suspension on the matter is permitted. But if strong level-connection obtains, it must be that we always have enough support for true beliefs about what rationality requires in different situations. That result may be hard to account for without supposing that we have a priori justification for true beliefs about the requirements of rationality.

EYAL TAL
BRANDEIS UNIVERSITY
EYALTAL@BRANDEIS.EDU

References

- Armstrong, D. M. (1963) 'Is Introspective Knowledge Incorrigible?' *Philosophical Review*, 72, 417–32.
- Basu, Rima. (2019) 'What We Epistemically Owe to Each Other'. *Philosophical Studies*, 176, 915–931.
- Bergmann, Michael. (2005) 'Defeaters and Higher-Level Requirements'. *Philosophical Quarterly*, 55, 419–436.
- Boghossian, Paul. (2008) 'Epistemic Rules'. *Journal of Philosophy*, 105, 472–500.
- Buchak, Lara. (2014) 'Rational Faith and Justified Belief'. In Laura Frances Callahan and Timothy O'Connor (eds.), *Religious Faith and Intellectual Virtue* (Oxford: Oxford University Press), 49–73.
- Chisholm, Roderick M. (1989) *Theory of Knowledge*. 3rd ed. Englewood Cliffs: Prentice-Hall.
- Christensen, David. (2007) 'Epistemology of Disagreement: The Good News'. *Philosophical Review*, 116, 187–217.
- Christensen, David. (2009). 'Disagreement as Evidence: The Epistemology of Controversy'. *Philosophy Compass*, 4, 756–67.
- Christensen, David. (2013). 'Epistemic Modesty Defended'. In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement: New Essays* (Oxford: Oxford University Press), 77–97.
- Christensen, David. (2016). 'Disagreement, Drugs, etc.: From Accuracy to Akrasia'. *Episteme*, 13, 397–422.
- Cohen, Stewart. (2013) 'A Defense of the (Almost) Equal Weight View'. In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement: New Essays* (Oxford: Oxford University Press), 98–120.
- Crawford, Sean. (2004) 'A Solution for Russellians to a Puzzle about Belief'. *Analysis*, 64, 223–29.
- Dorst, Kevin. (2020) 'Evidence: A Guide for the Uncertain'. *Philosophy and Phenomenological Research*, 100, 586–632.
- Dutant, Julien, and Clayton Littlejohn. (2016) 'Just Do It? When to Do What You Judge You Ought to Do'. *Synthese*, 195, 3855–772.
- Feldman, Richard. (2006) 'Epistemological Puzzles about Disagreement'. In Stephen Hetherington (ed.), *Epistemology Futures* (Oxford: Oxford University Press), 216–36.
- Feldman, Richard. (2007) 'Reasonable Religious Disagreements'. In Louise M. Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life* (Oxford: Oxford University Press), 194–214.
- Friedman, Jane. (2017) 'Why Suspend Judging?'. *Noûs*, 51, 302–26.
- Gardiner, Georgi. (2018) 'Evidentialism and Moral Encroachment'. In Kevin McCain (ed.), *Believing in Accordance with the Evidence: New Essays on Evidentialism* (New York: Springer), 169–95.
- Greco, Daniel. (2014) 'A Puzzle about Epistemic Akrasia'. *Philosophical Studies*, 167, 201–19.
- Hájek, Alan. (1998) 'Agnosticism Meets Bayesianism'. *Analysis*, 58, 199–206.
- Hazlett, Allan. (2012) 'Higher-Order Epistemic Attitudes and Intellectual Humility'. *Episteme*, 9, 205–23.
- Horowitz, Sophie. (2014) 'Epistemic Akrasia'. *Noûs*, 48, 718–44.
- Huemer, Michael. (2011) 'The Puzzle of Metacoherence'. *Philosophy and Phenomenological Research*, 82, 1–21.
- Jackson, Elizabeth. (2020) 'Belief, Credence, and Evidence'. *Synthese*, 197, 5073–92.
- Joyce, James M. (2010) 'A Defense of Imprecise Credences in Inference and Decision Making'. *Philosophical Perspectives*, 24, 281–323.
- Kiesewetter, Benjamin. (2016). 'You Ought to ϕ Only if You May Believe That You Ought to ϕ '. *Philosophical Quarterly*, 66, 760–82.
- Lasonen-Aarnio, Maria. (2020) 'Enkrasia or Evidentialism? Learning to Love Mismatch'. *Philosophical Studies*, 177, 597–632.
- Moss, Sarah. (2018) 'Moral Encroachment'. *Proceedings of the Aristotelian Society*, 118, 177–205.
- Raleigh, Thomas. (2021) 'Suspending is Believing'. *Synthese*, 198, 2449–74.

- Rosenkranz, Sven. (2007) 'Agnosticism as a Third Stance'. *Mind*, 116, 55–104.
- Schechter, Joshua. (2013) 'Rational Self-Doubt and the Failure of Closure'. *Philosophical Studies*, 163, 429–52.
- Schoenfeld, Miriam. (2015) 'A Dilemma for Calibrationism'. *Philosophy and Phenomenological Research*, 91, 425–55.
- Skipper, Mattias. (2021) 'Does Rationality Require Higher-Order Certainty'. *Synthese*, 198, 11561–85
- Smithies, Declan. (2012) 'Moore's Paradox and the Accessibility of Justification'. *Philosophy and Phenomenological Research*, 85, 273–300.
- Smithies, Declan. (2015). 'Ideal Rationality and Logical Omniscience'. *Synthese*, 192, 2769–793.
- Staffel, Julia. (2015) 'Beliefs, Buses, and Lotteries: Why Rational Belief Can't Be Stably High Credence'. *Philosophical Studies*, 173, 1721–34.
- Sturgeon, Scott. (2010) 'Confidence and Coarse-Grained Attitudes'. In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*. vol. 3 (Oxford: Oxford University Press), 126–49.
- Titelbaum, Michael. (2015) 'Rationality's Fixed Point (Or: In Defense of Right Reason)'. In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 5 (Oxford: Oxford University Press), 253–94.
- Turri, John (2012) 'A Puzzle about Withholding'. *Philosophical Quarterly*, 62, 355–64.
- Weatherston, Brian. (2019) *Normative Externalism*. Oxford: Oxford University Press.
- Worsnip, Alex. (2018) 'The Conflict of Evidence and Coherence'. *Philosophy and Phenomenological Research*, 96, 3–44.