Enkrasia or evidentialism? Learning to love mismatch

Maria Lasonen-Aarnio

Draft, please do not circulate

Abstract

I formulate a resilient paradox about epistemic rationality, discuss and reject various solutions, and sketch a way out. The paradox exemplifies a tension between a wide range of views of epistemic justification, on the one hand, and *enkratic requirements* on rationality, on the other. According to the enkratic requirements, certain mismatched doxastic states are irrational, such as believing *p*, while believing that it is irrational for one to believe *p.* I focus on an *evidentialist* view of justification on which a doxastic state regarding a proposition *p* is epistemically rational or justified just in case it tracks the degree to which one's evidence supports *p*. If it is possible to have certain kinds of misleading evidence (as I argue it is), then evidentialism and the enkratic requirements come into conflict. Yet, both have been defended as platitudinous. After discussing and rejecting three solutions, I argue that in order to solve the paradox, we should look to a more general phenomenon, namely, the way in which any rule or norm generates *derivative* rules or norms. I sketch the beginning of an account that rejects the enkratic requirements as genuine requirements of rationality, while nevertheless explaining their appeal.

I A Paradox

Consider a classic example first given by Bernard Williams: you are standing at a cocktail party, holding a glass that you believe, with good reason, to be gin. However, the glass in fact contains petrol. Conversely, assume that you are at a garage, and see a container you believe to be petrol, but that in fact contains gin. In both cases you very much desire a drink of gin. Do you have a reason to drink from the glass in the first case, and from the container in the second? A distinction between subjective and objective reasons is often invoked to answer this question: in the first case you have a subjective reason to take a sip, but have an objective reason not to, and in the second the converse holds. On one rough way of characterizing this distinction, whereas what a subject objectively ought to do can depend on factors she has no access to, what she subjectively ought to do depends on the information or evidence that she possesses.[1]

[1] Just how the distinction should be drawn is, of course, itself a source of dispute. For instance, instead of thinking of subjective reasons in terms of a subject's *evidence*, Schroeder (2009), for instance, draws the distinction in terms of a subject's *beliefs*. I do not wish to take a stance on exactly how the distinction here should be drawn, nor do my main points rely on it.

Justifying reasons for *belief* appear to fall squarely on the subjective side of this distinction, being restricted, in some sense, to the subject's perspective on the world.[2] Indeed, one might think that talk of *epistemic rationality* or *justification* just *is* talk of the subjective ought of epistemology: perhaps what one objectively ought to do is believe a proposition just in case it is true, but what one subjectively ought to believe depends on one's reasons or evidence. And so it may, at first sight, seem that epistemologists don't need a *further* distinction akin to that between subjective and objective reasons: it's not as if there are the subjective justifying reasons for belief, and in addition, the ultra-subjective ones.

But alas, matters are not this simple.[3] For facts about what it is rational for one to believe are themselves objective facts, and most epistemologists these days, whether of an internalist or externalist bend, would deny that we have full access to such facts. Just as there can be misleading evidence about whether a container contains gin or petrol, there can be misleading evidence about what it is rational or justified for one to believe. Indeed, a flourishing literature in epistemology is largely concerned with evidence bearing on just what one's evidence supports, or what it is rational to believe given one's evidence.[4] For starters, it isn't always clear just what one's evidence or reasons are. Consider any view on which having a proposition *p* as evidence requires bearing some (epistemic) relation to *p*. Call that relation R. In so far as bearing R to *p* doesn't entail bearing R to the proposition that one bears R to *p*, *p* might be part of one's evidence, even if it is not certain on the evidence that *p* is part of one's evidence. But even an RR-thesis wouldn't guarantee the result that it is always rational to be certain what the evidence is. The evidence must not only be certain, of each item of evidence, that it is part of the evidence, but it must also be certain of a certain set, that *that* is all the evidence there is. And even when it is certain just what the evidence is, it might not be certain what it is rational for one to believe given the evidence. Indeed, this seems to be assumed in much of the recent discussion of disagreement and defeat by higher-order evidence.

And so it looks like something reminiscent of the gin and petrol puzzle arises in epistemology. There isn't just the world and one's perspective on it. There is also one's perspective on one's perspective on the world (and so on). Assume, for instance, that *p* is likely on the subject's evidence. However, it is also likely on her evidence that *p* is not likely and hence, that it is not rational for her to believe *p*. What doxastic state regarding *p* ought the subject to have in such a situation?[5] What we have, in effect, is the makings of a paradox about epistemic rationality. The territory we are in is thorny, so in order to spell out the paradox, it will help to have a concrete view of epistemically rational or justified belief on the table. I will take as my starting point an *evidentialist* view. Indeed, I want to argue that there is a serious tension between such evidentialism, on the one hand, and what I will refer to as *enkratic requirements* on rationality, on the other. However, as I will try to make clear,

[2] For this talk of a subject's "perspective on the world" see, for instance, Alston (1985).

[3] Against, for instance, Feldman (1988). I sketch a more detailed argument below.

[4] I have in mind the literature on higher-order evidence – and a large part of the literature on disagreement can be counted in this camp.

[5] Not surprisingly, epistemologists sometimes invoke a distinction between subjective and objective (epistemic) obligations or justification. See, for instance, Alston (1985), Gibbons (2006), Goldman (1986: 73), and Pollock (1979).

evidentialism is not essential for creating trouble with these requirements, for problems arise given any of a wide range of views about epistemic rationality or justification.

Evidentialism is often characterized as a supervenience claim on which facts about what doxastic states it is rational for a subject to be in are fixed by her evidence.[6] But I, together with many others, find plausible the claim that the rational or justified doxastic states are those that *track degrees of evidential support*. I will be assuming that support is a probabilistic relation, though I don't think this follows from evidentialism as such. The idea, then, is that a doxastic state in a proposition p is epistemically permitted if and only if it tracks the probability of p on one's evidence, or the evidential probability of p.[7] A simple, somewhat natural way to think is that one is only permitted (and perhaps even *required*) to have credences that *match* evidential probabilities. So, for instance, if the probability of p is 0.9, one is permitted to assign to p a credence of 0.9, and no other credence. However, I want to leave the notion of tracking somewhat vague. For instance, the above characterization of evidentialism leaves open a view on which credences that are sufficiently *close* to evidential probabilities are permitted. So, for instance, if the probability of p is 0.9, one may be permitted to assign to p any credence sufficiently close to 0.9.

Now consider misleading evidence about what one's evidence supports, or about what one's evidential probabilities are. In so far as evidence can be misleading about itself, it would seem possible for it to be even *radically* misleading. Consider, then, the following kinds of evidential situations:

> p is likely on the evidence, but it is also likely that p is not likely on the evidence.
>
> p is not likely on the evidence, but it is likely that p is likely on the evidence.

---

[6] See, for instance, Feldman and Conee (2004), though note that I am not using 'evidentialism' to refer to the epistemological package they defend (including an internalist view of evidence). Note also that 'evidentialism' is sometimes used to refer to a view on which only evidential considerations can act as reasons for beliefs, or a view on which only such considerations can *rationalize* belief. Evidentialism as I intend it is just a claim about *epistemic* reasons; it is compatible with claiming that there can be practical reasons for belief.

[7] At first sight such a claim looks to entail, but not be entailed by, the supervenience claim. However, I intend my characterization of evidentialism to not even entail the supervenience claim. Here are two reasons why the entailment might fail. One issue has to do with how one thinks about evidential probabilities. Take, for instance, a view on which the evidential probability of a proposition (for a subject s) is given by conditionalizing the prior probability function that is ideally rational for s on s's total evidence. If different prior probability functions can be rational for different subjects (depending on features of their worlds, or perhaps even features of their cognitive faculties), then the tracking claim won't entail supervenience. (Indeed, what I say below is even compatible with a radical subjective Bayesian view on which *any* prior probability function is rational.) Second, I want to leave open a reading of 'tracking' on which what counts as tracking might depend on pragmatic factors such as stakes. So, for instance, in lower-stakes situations believing p might count as tracking one's evidence when the evidence makes p likely to degree 0.9, whereas in some higher-stakes situations the evidence must make p likely to at least to degree 0.99.

That evidence can have very poor access to itself isn't just an idle possibility that cannot completely be ruled out. Very few think that there is a viable account of evidence on which evidence always has *perfect* access to itself, and the kinds of reasons for thinking this support the idea that sometimes evidence might have *poor* access to itself. But more on this below.

The evidentialist has a simple answer regarding what doxastic states subjects in the kinds of evidential situations envisaged should (or are at least permitted to) adopt. Consider, for instance, situations in which *p* is in fact likely on the evidence, but it is likely that *p* is *not* likely on the evidence. In such a situation it seems that evidentialism would permit believing *p* while believing that the evidence does not support *p* and hence, that it is irrational for one to believe *p*. But such states are often brought up as paradigm examples of irrationality! A subject in such states is failing epistemically by her own lights, failing to abide by her own beliefs or opinions about what she is required or forbidden to believe. Rationality in general, it is often urged, just *is* doing what makes sense given one's perspective or point of view. No wonder it seems odd to make assertions like "I am rationally required to believe that it is raining, but it is not raining"; or "It is raining, but it is irrational for me to believe that". Coherence (or "metacoherence") requirements that prohibit the kinds of mismatched states described above – or at least seemingly similar states that involve some sort of failing by one's own lights – have been immensely popular across different areas of philosophy. Indeed, such requirements are simply often assumed, without argument, as premises for further theorizing.

Very roughly, *enkratic* requirements prohibit subjects from failing to comply with their judgments or beliefs about what they ought to do (or, in the case of practical action, from failing to form intentions that accord with such beliefs). What I call the *Negative Enkratic Principle* (EP-) below prohibits believing (or having high confidence) that one is forbidden (from the epistemic perspective) to be in a doxastic state D, while being in D. What I call the *Positive Enkratic Principle* (EP+) prohibits believing, or having high confidence, that one is required (from the epistemic perspective) to be in a state D, while failing to be in D. A subject who violates the enkratic principles either does something she takes to be forbidden from the perspective of epistemic rationality or justification, or fails to do something she takes to be required. If her beliefs about what she is forbidden or required to do were themselves irrational, we might simply require her to give up her higher-level beliefs. But these beliefs, it was assumed, were themselves arrived at by believing in accordance with her evidence. If evidence can be radically misled about what the evidence supports, and evidentialism is true, then it looks like subjects will sometimes be permitted (if not required) to violate the enkratic requirements.[8]

[8] As pointed out above, principles in the ballpark of EP+ and EP- have been very popular. For instance, Smith (1994: 178) argues that it would be irrational to believe that one's fully rational self would believe *p*, while failing to believe *p*. Wedgwood (2002) argues that it is a constitutive feature of the concept *rational belief* that judging that a given belief would not be rational commits one to not holding that belief. Kolodny (2005) formulates two core requirements of rationality that are very similar to EP- and EP+, arguing that they give rise to *all* other requirements of rationality. Huemer (2011) argues that it is irrational to believe a proposition *p* while believing, or even suspending judgment about whether one knows *p*. Gibbons (2006: 29) writes: "Surely there is something wrong with A-ing when you think you shouldn't. But there is also something wrong with failing to A when you should A." Christensen (2010b: 121) pretty much takes EP-, with "B" read as standing for belief, as a premise, though in numerous places

We now have a paradox. On an evidentialist view, a subject is permitted (perhaps even required) to be in states that track her evidential probabilities. But the question of which states track evidential probabilities is in itself an objective matter, and one that evidence can be misled about. For instance, even if a proposition *p* is likely on the evidence, it could be likely that *p* is not likely. Then, it looks like a subject could be permitted, perhaps even required, to believe that *p*, *and* that *p* is unlikely on her evidence – and hence, that given the evidential norm, she is forbidden to believe *p*. This is to say that subjects sometimes have an epistemic permission to be in akratic states. If there are also enkratic norms on epistemically permitted, justified, or rational doxastic states, something has to give. For how could perfectly proportioning one's doxastic states to the evidence force one to be epistemically *irrational*? If it cannot, then evidentialism, together with the possibility of evidence that is radically misled about itself, entails the falsity of the enkratic principles.

The assumptions that gave rise to the paradox were:

1. Evidentialism

2. Evidence can be radically misled about itself.

3. If 1. and 2., then the enkratic principles are false.

4. The enkratic principles are true.

Though my focus below will be on the paradox just sketched, it is worth emphasizing at the outset that an evidentialist theory of epistemically rational or justified belief is by no means essential for creating trouble with the enkratic requirements. The kind of paradox I have spelled out is surprisingly resilient.

Here is an abstract argument for why this is so. Take an epistemic theory that permits or requires one to be in a doxastic state just in case… (*fill in the dots with your preferred account*). Call the preferred theory Theory, and replace the first premise of the paradox by a premise stating that Theory is true. Now, for a wide range of candidates for Theory, there will be situations in which, for instance, Theory permits believing *p*, but also permits believing, or at least having high confidence, that one is not in the kinds of circumstances in which Theory permits believing *p*. That is, whatever the conditions for epistemically permitted doxastic

---

he defends the thought that "the rationality of first-order beliefs cannot in general be divorced from the rationality of certain second-order beliefs that bear on the epistemic status of those first-order beliefs" (Christensen 2007: 18). Smithies (2012) defends principles that yield versions of EP- and EP+. In fact, Smithies argues that it is not even rational to believe *p*, while suspending judgment on the question of whether believing *p* is rational. Titelbaum (2014) states that EP- (with "B" read as standing for belief, and "A" as standing for any doxastic state *or* intention) follows form the concept of rationality. For further defences of versions of EP-, see also Chisholm (1989: 6), Scanlon (1998: 25), Bergmann (2005), Elga (2005), Coates (2011), and Hazlett (2012).

states specified by Theory are, Theory sometimes permits false beliefs (or at least high degrees of confidence) about whether such conditions obtain. Replace the second premise above by this assumption. This is enough to get the paradox going. For instance, it now looks like Theory will sometimes permit being in a doxastic state (for instance, believing *p*), while permitting the belief that being in that state is impermissible by Theory. Indeed, as long as false beliefs can be permitted (or at least as long as one can be permitted to have high degrees of confidence in falsehoods), it would seem *prima facie* strange if Theory couldn't permit false beliefs about the subject matter of whether Theory permits a given doxastic state in one's circumstances.[9] And even if Theory does not permit false beliefs about its own application conditions, it might still permit false beliefs about which epistemological theory is correct in the first place. Indeed, in so far as epistemologists are in the business of forming epistemically justified or rational beliefs, it is difficult to see why false beliefs about matters epistemological could not nevertheless be rational.

To see just how resilient the paradox is, consider, for instance, a view on which it is necessary, but not sufficient, for believing *p* to be epistemically rational that *p* is (sufficiently) likely on one's evidence. In addition, it cannot be too likely that *p* is *not* sufficiently likely on one's evidence. Just to have some numbers on the table, assume the following condition:

> Believing *p* is permitted just in case *p* is sufficiently likely on one's evidence, and it is likely to *at most* degree 0.5 that *p* is not sufficiently likely

The worry now is that one's belief in *p* might be permitted by the new condition, even if it is likely on one's evidence that it is not. Assume, for instance, that *p* is sufficiently likely, and it is only likely to degree 0.3 that *p* is *not* sufficiently likely (and hence, likely to degree 0.7 that *p is* sufficiently likely). Nevertheless, one has misleading evidence about how likely it is that *p* is not sufficiently likely: in fact, it is very likely (say to degree 0.95) that it is likely that *p* is not sufficiently likely. Assume that the relevant subject knows the above condition. Then, in the situation described, it is likely on her evidence that she is not rationally permitted to believe *p*. For all that has been said, the belief *that she is not rationally permitted to believe p* can satisfy the above condition. In such a situation, even our new revised theory would permit the subject to both believe *p*, and to believe that it is rationally impermissible for her to believe *p*.

Hence, versions of the paradox spelled out above arise even when evidentialism is replaced by a range of alternative theories of epistemically rational, justified, or permitted belief.[10] However, in what follows I will focus on the paradox created by trying to combine evidentialism with the enkratic requirements. Some variant of an evidential norm strikes me as a very strong candidate for an overarching epistemic norm. But those not drawn to evidentialism at the outset can view the subsequent discussion as a kind of case study. I will argue that the kind of paradox spelled out above certainly shouldn't be taken as a reason to

[9] For a more detailed argument for a similar conclusion, see Lasonen-Aarnio (2014).

[10] As Alex Worsnip pointed out to me, a kind of coherentist theory on which all requirements of epistemic rationality take a wide-scope form like the enkratic principles to be discussed below, avoids the paradox. However, the vast majority of epistemologists would deny that there is nothing more to epistemic rationality than exemplifying a kind of coherence.

reject evidentialism, for the enkratic requirements actually stand on rather shaky ground. In this connection it is also worth emphasizing that several ways of defending certain formulations of enkratic requirements *assume*, rather than reject, an evidentialist condition on epistemically justified or rational belief. (Indeed, many recent proponents of enkratic requirements don't reject evidentialism.[11]) For instance, many have defended the idea that it is irrational for a subject who believes that her evidence does not support *p* to believe *p*. If such a view is motivated by the rough thought that rationally failing by one's own lights is, in itself, a failure of rationality, then it is being assumed that believing *p* in a situation in which one's evidence does not support *p is* a failure of (epistemic) rationality. But this, of course, assumes that (epistemic) rationality requires not holding beliefs that fail to be proportioned to the evidence. More generally: if something like evidentialism is not the correct story of epistemic reasons or epistemic rationality, then it is a prima facie mystery why a subject is required to avoid the kind of mismatch exhibited by, for instance, believing *p*, while believing that her evidence does not support *p*.

There is a range of rich, interesting responses to the paradox sketched above that I won't be able to discuss in detail. But let me briefly mention two broad views. The first responds to the paradox by accepting a pluralism of different notions of epistemically permitted or justified belief – basically, by denying that either of what was referred to above as objective and subjective justification is justification proper. There is the 'ought' (or 'may') of believing in accordance with the evidence, and the more subjective ought of believing in accordance with what the evidence says about the evidence (or one's justified beliefs about the matter). A subject with evidence that is misled about itself in ways described above is bound to violate one of these oughts. And two oughts wouldn't be enough to keep the problem at bay, for if evidence about evidence can be misleading, then evidence about evidence about evidence can me misleading, and so forth. Though I cannot give such a pluralist view the treatment it deserves, to my mind it seriously threatens the tradition giving the notion of epistemically permitted or justified belief a central role.[12] At any rate, before exploring other options, I don't think we should be quite so swift to accept a plurality of epistemic goods. In the end I want to argue that the paradox discussed above does not force us into such pluralism.

There is a second way of reconciling the evidential and enkratic norms that does not lead to a pluralism of different levels or orders of epistemic rationality. The idea is that though the two norms sometimes come into conflict, a doxastic state is epistemically permitted in the central, all-things-considered sense, only if it satisfies *both*. So, for instance, it is false that a subject is permitted to be in doxastic states that track her evidential probabilities, for sometimes such states violate the enkratic requirements. What, then, does the view under consideration say about situations in which evidence *is* radically misled about what it supports – what if, for instance, *p* is likely on the evidence, but it is also unlikely that *p* is likely? First, consider a construal of the evidential norm on which it never *requires* adopting doxastic states that track one's evidence, but merely permits such states. Then, the evidential norm would sometimes permit being in an akratic state, and the enkratic norm would forbid this. However, neither norm forbids adopting no doxastic state whatsoever regarding a proposition. So, for

[11] See, for instance, Horowitz (2014) and Titelbaum (2015).
[12] Though see Sepielli (2014) for a defence of a similar kind of pluralism.

instance, a subject with the kind of mismatched evidence described who simply adopts no opinion whatsoever regarding whether *p* wouldn't violate either norm. Now, one might worry that if the subject has considered the relevant question, then simply failing to adopt any opinion is not an option, and the proposed solution cannot get off the ground.[13] Besides, perhaps the evidential norm sometimes *requires* subjects to be in akratic states. Consider, then, a view on which situations in which evidence is radically misled about what it supports constitute *epistemic dilemmas*. In such situations the evidential norm will require a subject to be akratic, whereas the enkratic norm will forbid this. There is no way of satisfying both norms.[14]

If there is no other way of solving the paradox, we may yet have to resort to either variant of the kind of view just sketched, which constructs an account of epistemically permitted, justified, or rational belief by adopting both an evidential and enkratic norm. However, before seriously scrutinizing the case in favour of the enkratic requirements, I don't think we should be that swift to adopt enkratic norms as genuine requirements of rationality, thereby rejecting a purely evidentialist view.

I will spend the next sections discussing four remaining ways of resolving the paradox without rejecting evidentialism. The first rejects Premise 2 and hence, the assumption that evidence can be radically misled about itself. On such a view, if it is likely on the evidence that *p* is likely on the evidence, then *p is* likely on the evidence. I will outline a brief defense of Premise 2. I also say why rejecting it still leaves unresolved a similar paradox and hence, doesn't get to the heart of the problem. According to the second option, the problem is premise 3, which I shall term the *bridge premise*. According to this premise, the possibility of evidence that is radically misled about itself, together with evidentialism, entails that the enkratic principles are false. But there is a way of formulating the enkratic principles, the objection goes, on which the conflict disappears. There is a sense in which the third option, likewise, faults the bridge premise, and faults the way in which I have interpreted evidentialism and the enkratic principles. As is rather standard amongst epistemologists, I have spoken about epistemic justification and epistemic rationality in one breath. But one might object to the implicit assumption made that a theory of epistemic justification is a theory of *rationality*. This, one might hope, points to a way out of the paradox.

I will argue that all of these ways out of the paradox are problematic, and that we should opt for the fourth, which ultimately rejects the enkratic requirements. This raises the challenge for understanding why such requirements seem, at first sight, appealing. I argue that the right place to look is to a much more general phenomenon, namely, the way in which any norm appears to generate *derivative* norms that sometimes make recommendations that conflict with those of the original norm. The problem with epistemically akratic subjects is that while they might conform to the relevant primary epistemic norm, they violate a derivative norm. In this way, the kind of paradox spelled out above is, I want to urge, an instance of a much more general phenomenon. Offering a full account of the (seeming) normative force of derivative norms would take me well beyond the scope of this essay, but in the end I offer the beginning of such an account.

---

[13] Thanks to Alex Worsnip for pressing this point.

[14] This is one way of interpreting the kind of view that Christensen (2010a) defends.

I want to clear up a couple of potential misunderstandings, at the very outset, concerning the kind of view that someone who rejects the enkratic requirements is committed to. First, I am not defending a view on which evidence regarding whether one's evidence supports some proposition *p*, or evidence about whether it is rational to believe *p*, is *always* irrelevant for whether *p*. I am not saying anything like "how likely a first-order proposition *p* is on a body of evidence E depends on the subset of E that counts as first-order, and how likely higher-order propositions about how likely *p* is on the evidence depend on the subset of E that counts as higher-order". For one thing, I don't like the thought that the first- and higher-order components of a body of evidence can be separated; neither do I think that evidence bearing on higher-order questions is inert as far as first-order questions go. In general, that one has evidence supporting *p* is evidentially relevant for whether *p*, and vice versa. This can even be true in situations in which evidence is misleading regarding the question of whether it supports *p*.[15] The question of whether higher-order evidence has first-order relevance should be kept distinct from the question of whether such evidence can be misleading. Relatedly, I do not deny that there are no so-called higher-order defeaters. If I become confident that it will rain tomorrow based on evaluating a body of evidence, and someone whom I know to be good at evaluating such evidence tells me I have committed an error, this may well have the effect of lowering the probability that it will rain, even if I did in fact evaluate the evidence correctly.[16] Second, I am *not* claiming that any state involving epistemic akrasia is rational. A lot of the time the problem with akratic subjects is that they are failing to proportion their beliefs with their evidence. For instance, a subject knows that her evidence does not support the truth of astrology, while continuing to believe in astrology nevertheless, is irrational in virtue of failing to proportion her beliefs to her evidence. However, sometimes believing in accordance with the evidence forces one to be akratic, and it is these sorts of – perhaps somewhat special and unusual – situations that I want to focus on.

Before discussing how the paradox should be solved, it is worth taking a step back to discuss the enkratic requirements, as well as arguments in their favour, in a bit more detail.

II The enkratic requirements

As a very rough characterization, we can say that enkratic principles require that subjects comply with their own beliefs or judgments about what they ought to, or have reason to, do. Just as *practically* akratic subjects fail to act, or intend to act, in a way that reflects their judgments of what they have all-things-considered reason to do, *epistemically* akratic subjects

[15] Take, for instance, a case with the following structure. At a time t_0 the probability of a first-order proposition p is very low (say $1/30$) and the probability that at t_0 my evidence makes p likely is even lower (say $1/60$). At t_0 the probability of *both* propositions goes up: p is now $2/3$ likely, and it is $1/3$ likely that p is likely on one's my evidence. The end-result is a fairly akratic state, since I end up reasonably confident in p, but also reasonably confident that my evidence does not support p. However, the case cannot be faulted on the grounds that higher-order evidence is relevant for first-order questions, since *both* the probability of p *and* the probability that p is likely on my evidence goes up. Below I mention a case with this sort of structure.

[16] For reasons why we shouldn't think that higher-order defeat *always* occurs, see Lasonen-Aarnio (2013, 2014).

fail to be in doxastic states that reflect their judgments about what states they have epistemic reason to be in. What everyone seems to agree on is that the enkratic requirements are (at least putative) requirements of *rationality*. What is less clear is what the content of these requirements are. For now I will follow the somewhat standard practice of talking about epistemically justified or permitted belief, on the one hand – which, in a evidentialist framework, amounts to belief that tracks one's evidence – and epistemically *rational* belief, on the other, interchangeably. I will opt for a view that essentially characterizes epistemic akrasia as a mismatch between the doxastic states one is in, on the one hand, and one's beliefs (or states of confidence) about what doxastic states it would be epistemically rational for one to be in. One of the proposed solutions to the paradox, to be discussed below, contests this formulation. But to have something on the table, I will take the following principles as a starting point:

Negative Enkratic Principle (EP-)

Epistemic rationality requires that [if one believes (*or is confident*) that in one's current situation epistemic rationality requires **not** being in a state D, then one is **not** in D]

Positive Enkratic Principle (EP+)

Epistemic rationality requires that [if one believes (*or is confident*) that in one's current situation epistemic rationality requires being in a state D, then one is in D][17]

A permissibility-operator 'epistemic rationality permits that' stands to 'epistemic rationality requires that' as the possibility-operator in modal logic stands to the necessity-operator. 'D' stands for types of doxastic states, such as believing a proposition *p*, suspending judgment in *p*, disbelieving *p*, or assigning some credence *r* to *p*.[18] The enkratic principles are *wide-scope* requirements in the following sense: they prohibit being in a combination of doxastic states, but they don't, as such, prohibit or require one to be in a particular doxastic state, such as the state of believing *p*. For instance, EP- can be satisfied by *either* failing to believe (or be highly confident) that it is forbidden for one to be in S, *or* by not being in S. However, given the K-axiom of standard deontic logic, EP- and EP+ entail narrow-scope requirements. In effect, the resulting narrow-scope requirements state that rationally required belief (or high confidence) about what doxastic states one is forbidden or required to be in are factive.[19]

[17] *Assuming* that epistemic rationality is a matter of believing in accordance with one's epistemic reasons (according to evidentialism, proportioning one's doxastic states to one's evidence), we might as well replace talk of what epistemic rationality requires with what one's epistemic reasons require, or what one *ought*, in a distinctly epistemic sense of ought, to believe.

[18] Note that I am counting believing *p* and believing *q* as distinct types of doxastic states.

[19] Of course, such principles can be rejected on the grounds that rationality never requires one to believe or be confident in anything. But as was remarked above, the idea that subjects are required to proportion their doxastic states to their evidence is by no means foreign to epistemology. The K axiom is widely,

Recently epistemologists have defended, in particular, some version of the negative requirement EP-.[20] Couldn't they reject the positive requirement?[21] One might worry that such a view is poorly motivated, for numerous arguments and considerations in favor of EP- also favor EP+. Besides, a case can be made for thinking that adopting EP- within an evidentialist context more or less *entails* a cousin of EP+. Assume that at least sometimes subjects are positively required, and not merely permitted, to form beliefs that track their evidence. Now consider a subject who has evidence that permits her to be in an akratic state that violates EP+: it is likely to a sufficiently high degree on her evidence that she is required to be in a doxastic state D regarding a proposition *p*, but she is not in fact required to be in D. Assume, further, that the subject is *required* to believe that she is required to be in state D. Now, being required to be in D entails not being permitted to be in any alternative doxastic states D'. In so far as the subject has considered the question and recognizes the entailments in question, it looks like the subject is required to believe, of any alternative doxastic state D', that she is forbidden to be in D' (that is, requirements to believe are at least sometimes closed under logical entailment). But then, by EP-, no such D' *is* permitted.[22] The only candidate permitted state is D. We have derived the following narrow-scope principle: if one is required to believe that one is required to be in D, then one is required to either be in D, or to fail to adopt any doxastic state whatsoever regarding the relevant proposition.[23]

It is worth mentioning how *infallibilist* views of permitted belief can validate at least the versions of EP- and EP+ that deal with belief (and not merely high degrees of confidence). The majority of epistemologists assume, as almost platitudinous, that one can be permitted to believe falsehoods. But consider now a class of views that denies this. Assume that a subject is permitted to believe *p* only if *p* is entailed by her evidence, where evidence consists of true propositions. Or, consider views that impose a knowledge norm on belief on which a subject is only permitted to believe *p* if she *knows p*. As a result, beliefs about what doxastic states are forbidden, permitted, or required are only permitted if true. On such views, the versions of EP- and EP+ that deal with *belief* – as opposed to high *confidence* – come out as trivial corollaries of the factivity of permitted belief. As an example, consider a subject who fails EP- by being in a doxastic state D, while believing that she is forbidden to be in D. In order for the latter belief to be permitted, it must be true. But if it is true that she is forbidden to be

---

though not universally, accepted. Broome (2013: 120) rejects it; however, I am unpersuaded by his argument. But the main points I make below do not rest on whether or not the enkratic principles entail narrow-scope requirements.

[20] See Titelbaum (2014), Horowitz (2015).

[21] See, for instance, Goldman (2010), who argues that a subject might have justification to believe that she has justification to believe *p*, while lacking justification to believe *p*. Goldman at least leaves open the possibility that in such a situation a subject might be justified in disbelieving *p*, or suspending judgment in *p.*

[22] The reasoning employs the K-axiom. EP- requires one to satisfy a conditional. By the K-axiom this, together with the assumption that one is required to satisfy the antecedent of the conditional, entails that one is required to satisfy its consequent.

[23] One could argue, further, that if a subject has considered the question whether *p,* or if the proposition has been made salient to her, then adopting no doxastic attitude to *p* is simply not an option. And if the subject holds beliefs regarding the question of which doxastic states regarding *p* are required or permitted for her, then the proposition *p is* salient for her. Hence, in such situations failing to adopt any doxastic state whatsoever is simply not an option. Thanks to Alex Worsnip for discussion.

in D, then she cannot be permitted to be in D. Similarly, consider a subject who fails EP+ by believing that she is *required* to be in state D, but who fails to be in D. If her belief that she is required to be in D is to be permitted, it must be true. But then, she is forbidden to fail to be in D.

I am very sympathetic to infallibilism about permitted belief and hence, to the full belief versions of EP- and EP+. But the vast majority of philosophers drawn to enkratic principles do not defend them on these grounds. Besides, the kinds of infallibilist views sketched are not able to fully resolve the paradox, for while on such views permitted *belief* is factive, the same doesn't hold for permitted high degrees of confidence. After all, evidence consisting of truths can be misleading, supporting falsehoods to a high degree. As a result, these infallibilist views don't validate the enkratic principles formulated in terms of high confidence. But even the high-confidence versions of the enkratic principles have a lot of intuitive pull: for instance, a subject who believes *p*, while being *confident* that her evidence does not support *p* – and hence, that she is forbidden to believe *p* – seems to display a problematic sort of incoherence. But if akratic evidence is possible, then it looks like such akratic states can be permitted. Infallibilism about permitted belief still leaves this version of the paradox unresolved.

Many have found their favourite enkratic principles so plausible as not to be in need of further defense.[24] Those who defend such principles appeal to the intuition that there is something wrong with subjects who believe, assert, and act in accordance with akratic states. Consider, for instance, a subject who violates EP- by believing *p* while believing that her evidence does not support *p* and hence, that it is rationally forbidden for her to believe *p*. If a subject can be permitted to hold both of these beliefs, can't she at least sometimes also be permitted to believe their conjunction—and hence, to believe (1) or (1*) below? Similarly, it seems that subjects who violate EP+ could at least sometimes be permitted to believe – and, assuming one is permitted to assert propositions they are permitted to believe – to assert (2), (2*), or (3):

(1) *p*, but I am rationally forbidden to believe *p.*

(1*) *p*, but *p* is not likely on my evidence.

(2) ¬*p*, though rationality requires me to believe *p*.

(2*) ¬*p*, though *p* is likely on my evidence.

(3) *p,* but I am rationally required to suspend judgment on the matter.

[24] Mike Titelbaum (2014), for instance, thinks that EP- is conceptually true: "Just as part of the content of the concept *bachelor* makes it irrational to believe of a confirmed bachelor that he's married, the normative element in our concept of rationality makes it irrational to believe an attitude is rationally forbidden and still maintain that attitude" (p 289).

But such beliefs and assertions have a Moore-paradoxical feel to them.[25] A way of sharpening the seeming paradoxicality here further is by considering possible cases in which the evidence is radically misled about itself in the following way: it makes *p* likely, while making it likely that it makes ¬*p* likely.[26] Given evidentialism, it looks like a subject with such evidence could believe *p, but ¬p is likely on my evidence*. But this amounts to believing that one's evidence regarding *p* is misleading! It seems puzzling how a subject could ever rationally believe such a thing about her evidence. How can it be rational to believe *p* on the basis of evidence one takes to be misleading *as it bears on p*? Such a subject could, it seems, be in a position to rationally conclude that she got lucky in arriving at the truth, despite misleading evidence.[27]

Moreover, mismatched, akratic doxastic states seem to vindicate seemingly schizophrenic actions. For instance, Horowitz (2013: 11) observes that "it seems patently irrational to treat a bet about P and a bet about whether one's evidence supports P as completely separate". It would seem odd, for instance, to bet on a proposition *p* at 9:1 odds, while betting at 1:9 odds that those odds are irrational. The same kind of point, it seems, could be made against akratic suspension of judgment. Consider a subject who bets at 5:5 odds that it will rain within the next hour, but who bets only at 1:9 odds that her evidence supports rain to roughly degree 0.5. If you ask the subject how she can be sure that she interpreted her evidence correctly, it would be odd for her to say: "I am not at all sure that my evidence makes it rational to be 50% confident in rain. In fact, I am quite sure I shouldn't be offering such odds! But we are betting on whether it will rain and not on whether my evidence supports rain, right?"

I admit that the above kinds of arguments raise a challenge for those who deny the enkratic requirements. However, in the seeming costs of accepting the possibility of rational epistemic akrasia must be weighted with the costs of defending the enkratic requirements. The argument from Moore-paradoxicality assumes, first, that the best (and perhaps only) explanation of the seeming badness of believing propositions such as (1)-(3) is that such beliefs are rationally forbidden. As a result, the argument is only as strong as alternative explanations of this seeming badness are weak. Below I offer the beginning of an alternative line of explanation. And it is worth noting that infallibilists about permitted belief - for instance, those who defend a knowledge norm on belief and assertion - can accept that the relevant beliefs and assertions are rationally forbidden without having to accept the enkratic requirements as special requirements governing epistemic rationality. At the very least, then, arguments appealing to Moore paradoxicality must be coupled with a rejection of views that impose a factivity-entailing norm on belief and assertion.[28] Now, versions (albeit slightly weaker ones) of the challenge from Moore-paradoxicality arise even assuming an infallibilist view of permitted belief, for after all, infallibilism about permitted belief doesn't entail infallibilism about permitted high degrees of confidence. How, for instance, could it be

[25] See, for instance, Feldman (2005: 108), Bergmann (2005: 424), Hazlett (2012: 211), Huemer (2011), Smithies (2012).

[26] The latter does not follow from the former, but see Horowitz (2014) for a convincing example.

[27] See Horowitz (2014).

[28] Take, for instance, (1*). Assume that one is only permitted to believe or assert *p is not likely on my evidence* if this is true. If it is true, then assuming evidentialism, she isn't permitted to believe *p*. Similarly, if one is only ever permitted to believe truths, then one can never be permitted to believe conjunctions like *p, but my evidence supports ¬p* – and hence, that one's evidence regarding *p* is misleading.

rational for a subject to be *confident* that her evidence regarding a proposition *p* is misleading?[29]

But, as I argue in more detail below, evidentialists are hard pressed to deny that such states could ever be rational, for evidence can be radically mislead about what it supports – and, more generally, about what doxastic states are rational in one's present circumstances. Hence, anyone committed to evidentialism should be seriously worried about the enkratic requirements. (And, as I noted at the very outset, evidentialism is by no means essential for creating the paradox, for the enkratic requirements seem to conflict with a whole host of theories about epistemic rationality.) Indeed, the force of evidentialist considerations against the requirements is reflected in the fact that many who have recently defended some form of enkratic principle do, in fact, concede the existence of such cases, admitting that their favoured requirement fails in certain situations involving uncertainty about what one's evidence *is*.[30] But once it is admitted that the enkratic principles fail in certain situations, what of the above arguments, which seemed to support them across the board? Recognizing theoretical pressure to accept the existence of counterexamples doesn't make the seeming unpalatability of akratic states disappear. Hence, those who defend restricted versions of the principles face not only the challenge of providing some principled, informative way of circumscribing the problematic cases, but also of explaining away the intuition that certain kinds of mismatched states are always irrational (how, for instance, can it ever be rational to believe, or be confident, that one's evidence regarding *p* is misleading?).[31] In offering a framework in which the enkratic requirements have no useful work to do, and in attempting to explain away why so many have been drawn to them, I aim to make their failures more palatable. Moreover, I argue below that failures of the enkratic requirements are not restricted to cases involving uncertainty about what the evidence is, for evidence can also be misled about what the true support-facts are, as well as what the correct theory of epistemically rational or justified belief is in the first place.

At this point it is worth raising an immediate problem with EP- and EP+: it is not clear whether the kinds of motivations given for such principles apply in situations in which a subject lacks suitable access to her own mismatched states. First, consider a subject who believes *p*, but is rationally confident that she doesn't believe *p*, for it is likely on her evidence that she doesn't believe *p*. If the subject believes that she is forbidden to believe *p*, then she is failing to follow her own advice concerning what it is rational for her to believe, but there is a sense in which she is not irrational by her own lights, for she cannot access the mismatch within her mind. In light of such cases, proponents of enkratic requirements might want to restrict their preferred principles.[32] Note also that not any kind of access to one's own mind will do. Ama might believe *p*, and believe that Ama believes *p*, but nor realize that she is Ama and hence, not realize that *she* believes *p*. Or, she might realize that she is in a given doxastic

[29] Cf Horowitz (2014: 726).

[30] See, for instance, Horowitz (2014), Titelbaum (2015), and Elga (2005).

[31] Horowitz (2014) offers a promising way of meeting this challenge, though, as I argue below, failures of the enkratic requirements are not restricted to the kinds of cases she discusses.

[32] Cf. Titelbaum (2014), who defends EP-, but explicitly restricts the principle to cases that don't involve failures of "state luminosity". Broome (2010) also proposes restricting his preferred enkratic requirements.

state, without realizing that it is a state of *believing p*. Hence, modes of presentation matter. I suspect that controlling for such issues will prove to be tricky. If the enkratic requirements start crumbling in the face of such revisions, that only strengthens my overall case against them.

Given that the enkratic requirements are often defended on grounds of intuitiveness, it is worth flagging some further intuitions that conflict with them. First, consider situations of the following kind. You are sure that either *p* is likely or unlikely on your total evidence, but unfortunately, you have no idea which. In fact, you are roughly 50/50 about the matter. Perhaps, for instance, evidential experts are divided when it comes to the question, though they all agree that *p* is either likely or unlikely. Wouldn't it seem perfectly appropriate for you to suspend judgment about *p* in such a situation, adopting a mid-level credence in *p*? Many authors have expressed such verdicts. Here is one possible theoretical motivation: your *expectation* of the rational degree of confidence in *p* is mid-level, and rationality requires matching your degrees of confidence with your expectations about what is rational. But this would seem to put you into an akratic state, for you suspend judgment in *p*, despite being certain that suspension is not the evidentially supported attitude.[33] One attempt to avoid the problem restricts the enkratic requirements to the states of belief and disbelief. However, the problem still persists, at least if one thinks that there is some (perhaps context-sensitive) non-maximal credal threshold for belief.[34] Assume that in the context you are in, the threshold is 0.9, and you know this. You are 0.9 confident (and hence, believe) that the rational credence in *p* is 0.89, and 0.1 confident that it is 0.99. Now, if your credence in *p* ought to equal your *expectation* of the rational credence, then you ought to assign a credence of 0.9 to *p*. Hence, you ought to believe *p*, despite believing that it is irrational for you to believe *p*.[35] A more concessive response rejects the enkratic requirements in favour of some sort of *rational reflection* principle. I argue against rational reflection principles elsewhere.[36] For now let us focus on the enkratic requirements, for they have certainly had a lot of intuitive appeal.

Here is another consideration against an intuition-based methodology. A common defense of the irrationality of epistemic akrasia is appeal to the idea that failing by one's own lights is in itself a failure of rationality – recall the idea, mentioned above, that rationality just is doing what makes sense given one's perspective or point of view. However, not only does such an idea threaten the possibility of formulating any requirements of rationality, appeal to such an intuition has the potential to lead us into muddles, issuing conflicting verdicts about rationality. To see this, consider the seeming irrationality of going against one's judgments

---

[33] See also Lasonen-Aarnio (2015) for a discussion of the conflict between *rational reflection* principles and the enkratic requirements.

[34] Note that a view on which there is no systematic connection whatsoever between credence and belief is by no means guaranteed to avoid the problem!

[35] Cf Lasonen-Aarnio (2015). Another strategy would be to try to deny that the kinds of evidential situations described are possible. One might, for instance, suggest that the rational credences cannot be disjoint in the way assumed: If a subject assigns some non-zero credence to r_1 being the rational credence in *p*, and she also assigns some non-zero credence to r_2 being the rational credence, then she must assign a nonzero credence to *each* value between r_1 and r_2 being rational. Such a suggestion runs into obvious technical problems, for there are uncountably many reals between any two reals. Besides, more can be said to fill in the cases in which the kinds of disjoint distributions described do seem rational.

[36] Lasonen-Aarnio (2015)

about what one *overall* (and not just from the epistemic point of view) ought to do - and hence, the seeming irrationality of classic cases of akrasia. Assume that you recognize that your evidence does not support *p* (perhaps you recognize that it supports not-*p* instead) and hence, that you don't have epistemic reason to believe *p*. However, a *lot* would be gained by believing *p*. Perhaps, for instance, by some series of events, you know that your coming to believe *p* would save someone's life. Then, it might seem – at least to you – that your *overall* reasons require you to believe *p*. If you fail to believe *p*, you fail to adopt the doxastic state that, by your own lights, your *overall* reasons require you to adopt. But surely this is failing by your own standards! However, by believing *p*, you would be *epistemically* failing by your own standards. The intuition that failing by one's own standards is a failure of rationality gives us no clear verdict in such cases, for if anything, we have conflicting intuitions. Now, one might hope that the seeming conflict disappears as long as all the relevant principles are formulated as involving requirements that take wide-scope (as EP- and EP+). However, this is not obviously so. For assume (as many epistemologists do) that sometimes one is *required* to hold beliefs that reflect one's reasons or evidence. For instance, assume that you have considered all the relevant questions, and you recognize that it is likely on your evidence that you both have overall reason to believe *p* while lacking an epistemic reason to do so. Perhaps you are then *required* to believe that your overall reasons require you to believe *p*, while your epistemic reasons prohibit this. Then, it looks like the respective wide-scope requirements generate conflicting requirements: one requires you to believe *p* (since by not believing *p* you are being akratic), while the other requires you to not believe *p* (since by believing *p* you would be epistemically akratic).[37] No doubt there are ways of trying to avoid the conflict, but I think this shows that we shouldn't rely on mere intuitions to guide us through the rocky waters of rationality.

I will now go on to discuss four ways of resolving the paradox without giving up evidentialism. The first three attempt to reconcile evidentialism with the enkratic requirements. The first rejects premise 2. I have two complaints against this. First, such a view rules out as impossible evidential situations that are possible. Second, the paradox is really a manifestation of a more general problem that rejecting 2. does nothing to solve.

[37] Again, this reasoning relies on the K-axiom.

III Rejecting premise 2

I now want to briefly discuss a solution that rejects Premise 2. According to this solution, facts about how likely various propositions are on a body of evidence cannot be *too* unlikely on the evidence. Such a solution rules certain kinds of evidence to be metaphysically impossible: it is impossible to have evidence that permits believing, or being confident, in falsehoods regarding what attitudes are appropriate or permitted given the evidence. Arguments against such a view have been belabored by others. I will briefly canvass the kinds of considerations that persuade me that evidence can have even very poor access to facts about what it supports. Besides, rejecting Premise 2 still leaves open the possibility of having evidence that is misled about the true theory of what epistemic rationality consists of in the first place

First, I think there are compelling reasons to doubt that there is any view of evidence on which evidence is luminous in the following way: whenever some item *e* is part of one's evidence, it is certain on the evidence that *e* is part of one's evidence. If we think of evidence as propositional, and if there is some relation R that one must bear to a proposition *p* for *p* to be part of one's evidence, then one might bear R to *p* without bearing *R* to the proposition that one bears *R* to *p* (compare this with the *Positive Introspection Axiom* of modal logic). Pretty much any view of evidence is susceptible to a Williamsonian-style anti-luminosity argument.[38] But as was noted at the outset, even luminosity wouldn't guarantee the sort of access-condition on evidence that prevents the Paradox from arising. Assume that for every item of evidence I have, it is certain that I have that item of evidence. It still wouldn't follow that I am certain, of some total body of evidence E, that E is my evidence. In addition to luminosity, we need the condition that if an item *e* is *not* certain on my evidence, then it is certain on my evidence that *e* is not certain on my evidence. That is, if one *doesn't* bear R to *p*, one must bear R to the proposition that one doesn't bear R to *p*. (compare this with the *Negative Introspection Axiom* of modal logic).

Consider how both failures of positive and negative introspection create situations in which it is not certain on the evidence to what degree it supports various propositions. Assume that *p* is part of Mosi's evidence. In virtue of containing *p*, Mosi's evidence makes *q* likely – but *p* is the only item of evidence Mosi has that supports *q*. Because *q* is likely, Mosi is permitted to believe *q*, or at least be highly confident in *q*.[39] However, it is not certain on Mosi's evidence that *p* is part of his evidence. Perhaps, for instance, it is only 50% likely that *p* is part Mosi's evidence. Then, it may only be 50% likely on Mosi's evidence that his evidence makes *q* likely and hence, that he is permitted to believe *q*. Take a different case in which Mosi's evidence fails to negatively introspect itself. Proposition *p* is *not* part of Mosi's evidence, and neither does Mosi have any other item of evidence supporting *q*. Given how unlikely *q* is on his evidence, Mosi is permitted to disbelieve *q*, or at least to assign a very low credence to *q*. However, it is, say, 50% likely on Mosi's evidence that his evidence does contain *q*. As a result, it is 50% likely that Mosi is permitted (and perhaps even required) to be confident in *q*.

[38] See Williamson (2000, Ch 4).

[39] *q* may or may not be identical with *p*.

Someone might admit that evidence can fail to perfectly introspect itself, while still insisting that evidence cannot be *too* misleading regarding the question of what it is. It is very difficult to see what would guarantee such a result. Besides, there seem to be counterexamples available to even weak access conditions. I cannot here discuss such counterexamples, but the kinds of "clock belief" -types of cases that have been discussed in connection with rational reflection principles convince me that evidence can even be radically misled about what it consists in.[40] Indeed, these kinds of cases have persuaded numerous recent proponents of the enkratic requirements to restrict their favoured principles.[41] However, even if there was a motivated way of restricting the enkratic requirements so as to exclude such cases, I see no reason to think that they constitute the only kinds of counterexamples.[42]

Here is a very general line of argument. Let **F** be whatever property a doxastic state must have in order to be epistemically justified or rational. For at least numerous candidate properties **F**, a doxastic state D can have **F** even if another doxastic state – namely, the state of believing, or having high confidence, that D *lacks* **F** – itself has **F**. In fact, it is not at all easy to come up with a plausible candidate for **F** that rules out such a possibility.[43] If, in general, false beliefs (or at least states involving high degrees of confidence in falsehoods) can have **F**, why not false beliefs about what doxastic states have that property? Assume now that **F** is the property of tracking degrees of evidential support. It was noted above that there are two sources of uncertainty about one's evidence: first, uncertainty about what one's evidence *is*, and second, uncertainty about support-facts themselves – in a probabilistic context, uncertainty about how likely various propositions are on a determinate body of evidence E.[44] Radical rational uncertainty of either kind suffices to vindicate the second premise of the Paradox. Assume that one has a total body of evidence E, and it is certain on the evidence that E is one's total evidence. Still, if E is radically misled about what it supports, or what the evidential probabilities on E are, a given doxastic state might track the evidential probabilities, even if it is very uncertain on the evidence that it does.

Indeed, numerous authors have recently argued or assumed that rational uncertainty about what a determinate body of evidence supports is possible. Such an assumption is made, in particular, in connection with what is often referred to as "defeat by higher-order evidence". Here are the bare bones of a pretty typical kind of example of such putatively

---

[40] This is a term I borrow from Christensen (2010b), who discusses a kind of case first described by Williamson. See Lasonen-Aarnio (2015), Christensen (2010b), Elga (2013), Horowitz (2014), and Williamson (2014) for a discussion of such cases.

[41] See, for isntance, Horowitz (2014) and Titelbaum (2015).

[42] Horowitz (2014) offers a promising diagnosis of what is special about such clock belief -type cases.

[43] For a more detailed argument for a similar conclusion, see Lasonen-Aarnio (2014).

[44] The uncertainty involved in Williamson-type counterexamples to positive and negative introspection is of the first kind. Indeed, the kind of formal framework Williamson deploys for modeling such failures is not equipped to model cases in which the evidence certain about what it consists in, but uncertain about what it supports. This is because within the formalism, the evidence had by a subject (at a world) can be thought of as the set of accessible worlds, and the true facts about evidential support at a world are gotten by conditionalizing a unique prior probability function on this evidence. Then, there is a sense in which the formalism builds in certainty about facts about evidential support: it is always certain what the evidential support facts are, conditional on such-and-such being one's evidence. However, I take this to be nothing more than an idealization built into the formalism.

defeating evidence. At a time *t* a subject becomes confident in a proposition *p* based on evidence E. Her confidence is perfectly rational, for E makes *p* likely. However, she then acquires evidence that E does not support *p* to the degree she took it to. For instance, perhaps she acquires evidence that her ability to evaluate evidence has been compromised due to the effect of a drug, fatigue, or a condition like hypoxia; or perhaps she just hears that her competent friend became confident that *p* is false based on E. The higher-order evidence need not call into question what the original evidence was, but rather, whether it really does make *p* likely. It is urged that the subject can no longer be certain, or even reasonably confident, that *p* was likely on the original evidence E. The enkratic requirements are then often invoked to argue that the higher-order evidence acts as a defeater: rationality calls for lowering one's confidence in *p* in response to the higher-order evidence, for one would otherwise be epistemically akratic.[45]

However, I am worried that the resulting position is unstable, for the assumption that there can be rational uncertainty of even a rather radical sort about true support-facts can be turned into an argument *against* the enkratic requirements. Now, proponents of defeat and enkratic requirements can reply that the kind of case described need not involve uncertainty about what one's *present* total evidence supports, but rather, uncertainty about what the evidence one had at a slightly earlier time supported, a time prior to acquiring the relevant higher-order evidence concerning drugs hypoxia, or disagreement. However, once rational uncertainty about true support-facts is admitted, it seems *ad hoc* to maintain that such uncertainty could not concern one's present total evidence. For instance, if testimonial evidence from a reputable source can make it rational to doubt whether one's previous evidence really supported *p* (even if it did), why cannot such testimonial evidence make it rational to doubt whether one's *present* evidence supports *p* (even if it does)?[46]

Someone might still resist the idea that there could be misleading evidence about how likely a proposition *p* is on a determinate body of evidence E, for facts about evidential support are *a priori*. Titelbaum (2014), for instance, argues that rationality requires one to have certainty in truths about what various determinate bodies of evidence support. To explain why there couldn't be misleading evidence regarding such matters, he argues that because support-facts are *a priori*, the evidence for them must always be in place. But having evidence that in some situations (absent certain other evidence) lends support to a proposition *p* doesn't entail having *overall* evidence that supports *p*: the *a priori* knowability of a proposition does not entail that all subjects are always in a position to know it *a priori*. To say the least, it is at least a *prima facie* mystery why there couldn't be misleading evidence about what various determinate bodies of evidence support. It is difficult to see what else than a

[45] The argument assumes that the higher-order evidence makes it rational to have lower confidence not only in the claim that one's *previous* evidence E supported *p* and hence, that it was rational for one to believe *p*, but also in the claim that one's present total evidence supports *p* and hence, that it is rational for one to believe *p* in one's present situation. Only then would failing to lower one's confidence in *p* lead into a state of epistemic akrasia.

[46] Titebaum (2015) argues, in more detail, that such a view would be *ad hoc*. In particular, he argues that any adequate answer to the question of why the "justificatory map is arranged such that one is never all-things-considered justified in both an attitude *A* and the belief that *A* is rationally forbidden in one's current situation" will prevent not only false but rational beliefs about what is rational in one's *present* situation, but also false but rational beliefs about what is rational in *other* situations.

prior commitment to the enkratic requirements might entice one to think this – and indeed, Titelbaum's argumentation assumes EP- as a premise.

I have said why I think Premise 2 of the paradox is true. But there is a further, perhaps even deeper, reason to think that rejecting Premise 2 cannot be the way to go here, for it fails to solve a more general version of the paradox involving not uncertainty about one's evidence, but uncertainty about what the correct story of epistemic justification is in the first place. (What I say in connection with my positive view will cast further doubt on solutions that reject Premise 2, for there is yet another dimension along which the solution is not general enough.) Consider the following examples:

> *Moral uncertainty*
>
> Assume that the correct theory of objective moral rightness is some form of utilitarianism that urges performing the action that maximizes happiness. Assume that in your present circumstances, action a_1 maximizes happiness. Consider whether you ought to ϕ in the following kinds of situations: (i) you have excellent evidence that action a_2, not action a_1, maximizes happiness; (ii) you have excellent evidence that the kind of utilitarianism just sketched is incorrect and that instead, some deontological theory is correct that tells you to perform an altogether different action a_3.

> *Epistemic uncertainty*
>
> Assume that evidentialism is correct. Consider whether you ought to believe *p* in the following kinds of situations: (i) *p* is likely on your evidence, but it is unlikely on your evidence that *p* is likely on your evidence; (ii) you have excellent evidence for a non-evidentialist epistemological theory on which you are required to believe *p* in your current situation, even though your evidence does not support *p*.

The notion of subjective rightness, or subjective moral oughts, has often been invoked in connection with situations of the kind described in (i) of *Moral uncertainty*: one subjectively ought to perform the action that, in some sense, maximizes happiness from one's perspective. This might be performing the action that is most likely, given one's evidence, to be the objectively right, happiness-maximizing action, or it might be the action that yields the greatest *expected* happiness of all available actions.[47] But uncertainty about what actions are right might run even deeper. A subject might be uncertain what the correct account of the subjective ought is. Or, as in situation (ii), she might be uncertain what the correct first-order normative theory is in the first place.[48] A subject who fails to perform (or intend to perform) action a_3, while being highly confident that that is the morally right action to perform, could be characterized as morally akratic. But there doesn't seem to be any stopping to this process

[47] Jackson (1991) discusses and criticizes the first proposal, defending the second – though he formulates these views in terms of moral utility, not happiness.

[48] For a discussion of such cases see Ross (2006), Sepielli (2009, 2014), Lockhart (2000).

of generating ever more subjective moral oughts, for the subject might also be uncertain what the correct way of factoring in uncertainty about the correct first-order moral theories is.[49]

In *Epistemological uncertainty* (i) is the kind of situation that gave rise to our original paradox. The solution under discussion was to deny that it is possible for evidence to be radically misled about what it supports. But there is a way in which the uncertainty in (ii) runs even deeper. If the subject fails to believe *p*, she acts against her own judgments about what she is required, from the epistemic perspective, to believe. It is difficult to distinguish the sense that there is something irrational about such a subject from the sense that there is something irrational about a subject who believes *p*, despite believing that her evidence doesn't support *p*. In fact, it is not obvious whether there is anything incoherent about the latter subject *unless* she at least implicitly recognizes that she is not permitted to hold beliefs that go against what her evidence supports. This is one reason to formulate the enkratic requirements as I have, in terms of opinions about what it is epistemically rational for one to believe, rather than in terms of opinions about what the evidence supports.[50]

Hence, the case *Epistemic uncertainty* (ii) creates a paradox very similar to the original one: by evidentialism, the subject ought to proportion her beliefs to her evidence, but if she does so, she appears to be epistemically akratic by believing in a way that she, by her own lights concerning what her epistemic reasons require her to believe, ought not to believe. Again, the new paradox does not essentially rely on evidentialism. Whatever the true justification-making property **F** is, it seems that one could be in circumstances in which a false belief (or at least high degree of confidence) concerning what the true justification-making property is could itself have **F**. In so far as **F** is the justification-conferring property, it follows that epistemically akratic states could sometimes be rational. Ruling out evidence that is radically misled about what it supports does nothing to solve the resulting paradox.

The challenge is particularly pressing for what might be termed *content-neutral* theories of epistemic justification, theories on which the conditions for epistemically justified or rational belief don't depend on the content of the belief in question. Such a theory would not, for instance, allow for sufficient evidential support to be a sufficient condition for the justification of "first-order" beliefs – beliefs not concerning the justificatory status of candidate doxastic states – but not for beliefs concerning such "higher-order" matters,

---

[49] Cf. Sepielli (2014).

[50] Here is a toy theory. Justified (or epistemically permitted or epistemically rational) beliefs are beliefs that manifest certain kinds of belief-forming dispositions. Which ones? Ones that do a good job at yielding knowledge, understood as belief that could not easily have been false. Moreover, at least in certain cases the value of knowledge is so great that it is worth taking a shot at acquiring knowledge even if there is a considerable risk of forming a false belief. Take, in particular, a case in which there is a large set of candidate hypotheses concerning some issue (perhaps they are scientific theories), each of which has a low probability on the evidence, but one of which is guaranteed to be true. If I had some principled way of deciding which of these theories to believe, should the theory I believe turn out to be true, my belief could not easily have been false and hence, would constitute knowledge. So the view sketched might recommend believing a hypothesis even if there is little evidence that it is true. By no means am I advocating such a theory; the claim is merely that it might be epistemically permissible (or rational, or justified) for a subject to be confident that it is true. The toy theory is inspired by work by Jeremy Lent.

requiring instead a condition entailing that such higher-order beliefs are only justified if true.[51]

So here is where we are. I considered a view that rejects Premise 2 of the paradox, ruling out evidence that is radically misleading about itself as metaphysically impossible, thereby ruling out situations in which one holds false, evidentially supported beliefs about what doxastic states are required or forbidden as a result of holding false beliefs about what one's evidence supports. However, false beliefs about what doxastic states are required or forbidden need not rest on false beliefs about one's evidence, but on a false, non-evidentialist epistemological theory instead.

IV Rejecting premise 3 and re-formulating the enkratic requirements

The reasoning behind Premise 3 of the Paradox was very simple. As an example, take a case of akratic evidence in which a proposition *p* is likely on a subject's total body of evidence, but it is also likely on her evidence that *p* is not likely. According to Premise 2, such evidence is possible. At least if the subject knows that she is not permitted to believe (or be confident) in propositions that are not likely on her evidence, then according to evidentialism, she is then permitted to believe *p* (or at least permitted to assign high confidence to *p*), while being permitted to believe (or at least be confident) that she is not permitted to believe *p*. Hence, it is possible that a subject is permitted to believe *p* and believe that she is forbidden to believe *p*. We have a violation of EP-.

Evidentialism and the possibility of evidence that is misled about what it supports (that is, Premise 1 and Premise 2 of the Paradox) entail that it is possible to be permitted to be in an epistemically akratic state. The objection now is that such a claim is compatible with there being no possible subjects who are *permissibly* in akratic states. That is, upon coming to be in an akratic state, a subject's permission to be in that state might automatically disappear. Further, the objection goes, the enkratic requirements should be re-formulated as saying that there are no subjects who are permissibly in akratic states. Once the enkratic requirements are formulated correctly, the paradox dissipates, for Premise 3 is rendered false.

The proposed resolution of the paradox assumes that something like the following is true of *any* case in which a subject's evidence is akratic: she is permitted to be in an epistemically akratic state, but she is in no position to come to permissibly be in such a state.

---

[51] Littlejohn (*forthcoming*) seems to defend a view along these lines, thereby rejecting content-neutrality. His justification for imposing such a factivity-condition on the relevant higher-order beliefs is roughly the following. Assume, for instance, that the epistemic rationality of first-order beliefs consists in being sufficiently likely on the evidence. If one then falsely believes that this is not what the epistemic rationality of these beliefs consists in, one makes *the same kind of mistake* as when holding first-order beliefs that fail to be likely on the evidence. I cannot do justice to Littlejohn's view here, but I am not at all convinced that the mistakes in question are the same kind. Note also that it is difficult to apply Littlejohn's justification to cases in which there is no rational uncertainty about what epistemic justification or rationality involves in the first place, but there *is* rational uncertainty about whether the conditions imposed by what one knows to be the correct theory of epistemic justification in fact obtain. Moreover, Littlejohn's view still allows for states that many have deemed to be epistemically akratic: for instance, believing, or being confident, that one's evidence strongly supports a certain theory, while failing to believe, or be confident, in that theory.

But why? Because, one might suggest, coming to be in such an overall state necessarily generates a defeater for one of its component states. Take, for instance, a situation in which *p* is likely on my evidence, but it is also likely that *p* is not likely. Such situations, the thought goes, are possible: the mere fact that it is likely that *p* is not likely need not defeat evidence that, on its own, makes *p* likely. However, coming to *believe* (or be confident) that *p* is not likely constitutes such a defeater. It is simply not possible for one's evidence to make *p* likely in the presence of such a belief.[52] Though proponents of this strategy might want to defend only EP- but not EP+, one could argue, similarly, that believing that one's evidence requires one to be in a state D acts as a defeater for any doxastic *alternative* to D.

First, it is not at all clear that the enkratic requirements as I have formulated them, read as concerning what might turn out to be mere permissions, are that toothless. Consider a view on which a subject is permitted to believe that it is raining and believe that it is not raining, even if she can never permissibly hold both beliefs. (Perhaps, for instance, these beliefs are mutually defeating). The claim that there are no possible subjects who permissibly hold contradictory beliefs doesn't remove the implausibility of claiming that a subject can be *permitted* to holds such beliefs. Instead, we should insist that there is simply no epistemic situation that permits holding contradictory beliefs, since there is no situation in which that would be the epistemically appropriate thing to do. Similarly, proponents of enkratic requirements could dig in their heels here: there is no epistemic situation that permits being irrational by one's own lights.

Second and more importantly, I don't see how to run the proposal within an evidentialist framework. It would have to be assumed that merely forming a higher-order belief (or state of confidence) concerning what one is forbidden or required to believe always changes one's evidence in such a way as to act as the desired defeater. For instance, merely coming to believe that my evidence doesn't support *p* would change my evidence in such a way that it no longer supports *p* (though it did before I formed the higher-order belief). Unless we countenance false evidence, the evidence thereby acquired cannot be that my evidence doesn't support *p*. And unless one's mental states are luminous, the evidence acquired cannot be that I believe that my evidence doesn't support *p*. Perhaps the idea is that the mental state of believing that I am forbidden to believe *p* itself comes to be part of my evidence. But a story still needs to be told of why adding such a mental state to one's evidence *always*, irrespective of what other evidence one has, has the result that the total evidence no longer supports *p*. It is far from clear what such a story might look like. Further, the present proposal is threatened by counterexamples. Just like a view that rejects Premise 2 of the paradox, it makes a claim about metaphysically possible evidence. The claim isn't that akratic evidence is impossible, but that akratic evidence is impossible *when one is in certain doxastic states*. For instance, it is impossible for a subject who believes, or is confident, that she is forbidden to believe *p* to have evidence that supports *p*. But such restrictions on evidence strike me as implausible for reasons discussed in connection with the solution that rejects Premise 2. Indeed, the case of *clock beliefs* mentioned above also constitutes a counterexample to the present view.

[52] Bergmann (2005) defends a view that comes very close to this strategy: he argues that the belief that one's belief in *p* is not reliably formed – whether or not this belief itself is justified – defeats whatever justification one may otherwise have had for *p*.

At least in some situations involving uncertainty about one's evidence, there seem to be no rational alternatives to akratic states.

Finally, the defeat-strategy seems hopeless in ruling out possible cases in which a subject has mismatched evidence making it rational to suspend judgment in a proposition *p*, despite rationally believing that it is irrational for her to do so, since her evidence makes *p* likely.[53] In order to apply the strategy to such akratic states, it would have to be argued that believing that one's evidence makes *p* likely acts as a defeater for the state of suspending judgment *by making it rational to believe (or be confident in) p*. But this would seem to commit one to the claim that merely coming to believe that *p* is likely on one's evidence changes one's evidence so that *p* now *is* likely on one's evidence! Surely it is not that easy to acquire evidence for any claim whatsoever.[54]

V Rejecting Premise 3, take 2: The ought of reasons and the ought of rationality

At the very outset I mentioned a way out of the paradox that countenances a plurality of different notions of epistemic justification. But there is a different kind of pluralism that has yet to be explored, one that is much more friendly to evidentialism. The pluralism in question draws on a distinction between the '*ought' of reasons* and the '*ought*' *of rationality*.[55] *Rationality* concerns, roughly, a kind of internal coherence amongst a subject's attitudes. And rationality requires, for instance, that a subject not manifest the kind of incoherence involved in being epistemically akratic. By contrast, what is permitted (and perhaps required) given one's epistemic *reasons* is proportioning one's doxastic states to the evidence. So whereas evidentialism deals with the 'ought' or 'may' of reasons, the enkratic requirements deal with the 'ought' or 'may' of rationality. One's epistemic reasons may permit, and perhaps even require, being in an akratic state, but rationality never permits this.[56]

It is important to see that the present solution does not in any way follow from admitting a distinction between the 'ought' of epistemic reasons and the 'ought' of rationality. In particular, it would not suffice to argue that evidentialism should not be thought of as a norm of rationality, and hence, that *rationality* doesn't require proportioning one's beliefs to the evidence. For even if being rational does not entail believing in accordance with one's epistemic reasons, one might still think that the converse holds.[57] Indeed, many would deny that a subject's doxastic states can perfectly accord with her epistemic reasons in being not only propositionally, but also doxastically, justified, while the subject still flouts some norm of rationality. That epistemic requirements could sometimes conflict with, say, the

---

[53] Worsnip (ms.) also makes this point.

[54] Not surprisingly, while Bergmann (2005) argues that merely having a certain belief can act as a defeater (whether or not the belief in question is justified), he thinks that "unjustified supporting beliefs cannot confer justification" (p. 426). Those sympathetic to Bergmann's thought here would have to combine the defeat solution with denying that the enkratic requirements hold for suspension of judgment. However, it is difficult to prevent the kinds of positive arguments and motivations for the enkratic requirements from applying to at least some cases involving suspension: for instance, a subject who suspends judgment in *p*, despite believing that it is irrational for her to do so, since her evidence makes *p* likely, appears incoherent in just the same way as a subject who believes *p*, despite believing that it is irrational for her to do so, since her evidence does not make *p* likely.

requirements of morality is not surprising; but the conclusion that a subject who fully satisfies such requirements could nevertheless be irrational is much less palatable. Consider paradigm coherence-norms on belief, such as "don't hold contradictory beliefs!" or "believe (obvious) entailments of what you believe!". A subject who perfectly proportions her beliefs to her evidence never violates such norms. Evidence can never support contradictions. And if the evidence supports *p*, and *p* entails *q*, then (at least if the entailment is sufficiently obvious), the evidence also supports *q*. Indeed, one of the big advantages of an evidence-norm on belief is that conforming to it guarantees conforming to various coherence-norms!

We cannot simply look to ordinary talk as a guide to the notion of rationality required by the present solution to the paradox, for it is completely natural to describe a subject who ignores part of her evidence, or whose evaluation of her evidence is influenced by her desires, as irrational. Proponents of the distinction between reasons and rationality often characterize rationality as a matter of something like internal coherence within a subject's mind – perhaps, for instance, of a specific kind of coherence among her non-factive attitudes.[58] But this in itself does not suffice to exclude evidential norms as requirements of rationality, for one could argue that subjects who fail to proportion their beliefs to their evidence *do* manifest a kind of incoherence.[59] Now, perhaps considering certain combinations of attitudes from a first-person perspective will help here. Alex Worsnip emphasizes that certain combinations of attitudes – for instance, believing *p* while believing that one's evidence does not support *p* – are not "capable of withstanding serious reflection" from a first-person perspective.[60] Believing something, he notes, is taking it to be true. But then, it is difficult to make sense of a state of believing a proposition *p*, while believing that *p* isn't at all likely to be true in light of all the available information.

None of this is to deny that at least in a lot of cases, it *is* difficult to make sense of certain combinations of attitudes. Below I argue that we can explain just why they seem so perplexing without having to commit ourselves to the idea that there is a *sui generis* kind of normativity that rules them out. My main worry with the idea that certain states are difficult to make sense of from a first-person perspective is that it is too slippery to adequately delineate

---

[55] For such a distinction see, for instance, Kolodly (2005: 509-510), Broome (2013), Scanlon (1998, Ch 1), Davidson (1985). The idea that the 'oughts' of reasons and rationality can conflict is not new. Such conflicts are discussed, for instance, by Kolodny (2005).

[56] I am grateful to excellent work by Alex Worsnip (ms.) for drawing my attention to this option out of the paradox.

[57] Now, it is true that a subject might have all and only the doxastic states her epistemic reasons permit or require her to have, while still failing for the reason that her doxastic states are not based on the right reasons. That is, she might be in a state just in case it is propositionally justified for her, even if some of her doxastic states fail to be doxastically justified. However, proponents of the present strategy would hardly admit that such a failing is a failing of *rationality*. I am indebted to a discussion with Han van Wietmarschen.

[58] See, for instance, Broome (2013: 152).

[59] Consider, for instance, a theory of evidence on which a subject's perceptual states constitute evidence. Now consider a subject undergoing a paradigm perceptual experience as of rain, who has no reason to distrust her experience, but who fails to believe that it is raining. It would not be at all unnatural to say that there is incoherence within the subject's mind, for her beliefs don't match her experiences.

[60] See Worsip (ms.). Worsnip argues that such states are even more difficult to make sense of than states of practical akrasia.

the intended notion of rationality. In particular, it always seems possible to fill in the details of a given a case so as to make better first-person sense of an epistemically akratic combination of attitudes. Recall some of the points made above, having to do with poor access to one's own mind, or access involving the wrong modes of presentation: it is not that difficult to make sense of a subject who, for instance, believes *p*, believes that *p* is highly unlikely on her evidence, but who also believes that she *doesn't* believe *p*. Or, consider subjects who are in the grips of false theories of epistemic justification. Take a subject who believes, in response to evidence she has, that in some situations her epistemic reasons require her to believe against her evidence – when, for instance, this is her best shot at acquiring knowledge.[61] Or take a subject who manifests a more traditional kind of incoherence by believing *p* and believing not-*p*, while also believing herself to be in an evidential situation that requires her to hold such contradictory beliefs, since they are supported by her evidence. Such considerations show the limits of a defense of the enkratic requirements appealing to the idea that "failing by one's own lights is in itself a failure of rationality". As Broome (2013: 91-93) notes, this idea makes it impossible to formulate any general requirements of rationality. But if the difficulty of making sense of epistemically akratic states from a first-person perspective does not result from the fact that at least paradigmatically, such states involve failing by one's own lights, what, then, is the difficulty supposed to consist in?

The claim that a subject who responds perfectly to her epistemic reasons is nevertheless irrational is surprising, and there are no clear precedents of such situations.[62] Drawing a distinction between norms of reason and norms of rationality is far from equivalent to any such claim. Indeed, my main reservation with the present solution has to do with its commitment to a notion of rationality that requires *more* than the kinds of widely accepted coherence requirements the satisfaction of which is guaranteed by perfectly proportioning one's doxastic states to the evidence ("don't hold contradictory beliefs"; "make your credences conform to the axioms of probability!"), but *less* than actually proportioning one's beliefs to the evidence. However, it is worth emphasizing that in a certain respect the present suggestion is close to the kind of account I want, in the end, to defend, for it concedes that the normativity (or, perhaps, apparent normativity) of the enkratic requirements is not the normativity of epistemic justification. The strategy I defend below, however, is more minimal in avoiding a commitment to a *sui generis* kind of 'ought' attaching to the enkratic requirements. Prima facie at least, there are reasons to prefer such an account: not only is it more parsimonious, but it avoids tricky questions having to do with how it is that we can be governed by such conflicting, *sui generis* 'oughts'.[63]

---

[61] See footnote 50.

[62] At this point someone might appeal to Moore's classic paradox: consider the Moorean proposition *p, but I don't believe p*. Such a proposition could be likely on one's evidence, but it would seem irrational to ever believe it. I agree that Moore's paradox is highly analogous to the present paradox, and a lot of the moves one can make to solve one paradox carry over to the other. However, for that very reason, I don't find appeal to Moore's paradox to back up the solution under discussion to be dialectically effective.

[63] Cf. Kolodny (2005: 555-557).

VI Rejecting the enkratic requirements (rejecting premise 4)

We are still left with the paradox. However, it would be too swift to reject evidentialism on this basis. I argued that the paradox is surprisingly resilient, and arises for a variety of alternative theories of epistemically justified or rational belief. Besides, it is far from obvious that the enkratic requirements are true. They should not be adopted on the basis of intuitiveness alone, for it is not difficult to generate intuition-based verdicts that conflict with them.

The arguments put forth so far constitute a negative case in favour of a solution to the paradox that rejects the enkratic requirements. But before concluding, I want to say something more positive about what the source of the seeming normative force of these requirements is. What follows is a promissory note concerning how I think the paradox ought to be solved. In particular, it arises, I want to urge, from a general puzzle concerning the way in which any norm seems to generate secondary or derivative norms. We need to give an account of this more general phenomenon, and whatever the correct account is, it will explain our sense that there is something wrong with a subject who adopts doxastic states that go against her own views about what states are epistemically rational, or what is supported by her evidence.

Consider any rule or norm that takes roughly the following form:

N_0 ϕ just in case you are in circumstances C!

The rule might pertain to morality, epistemology, or to mundane matters such as making good green tea ("pour the water into the cup just in case it is 90°C!"). Now consider situations in which your evidence is highly misleading regarding whether or not you are in C: Perhaps you are in C, but it is highly likely on your evidence that you are not. Or, perhaps you are not in C, but it is highly likely on your evidence that you are. As a simple example, assume that you are trying to make good green tea, and have done everything that can be reasonably expected to ensure that the water in the kettle is 90°C. However, your evidence is highly misleading, for in fact, the water hasn't heated up at all. Given your situation, *should* you pour the water into the cup? Or, are you at least seemingly *permitted* to do so? Here is the observation I want to begin with: if you pour the water into the cup in such circumstances, we feel pressure to evaluate your action positively. Indeed, it even feels natural to say that you *ought* to pour in the water, and that in some way, it would be *irrational* for you not to do so.

This raises an explanatory challenge: clearly, if you have highly misleading evidence to the effect that you are in circumstances C, our original norm N_0 itself does not urge you to ϕ, since you are not in C. According to N_0, it is *being* in C, and not *believing* that you are C (even if in response to good evidence), that gives you a reason to ϕ. Nevertheless, there is a seeming ought to ϕ when all your evidence points to being in C, an ought that in some sense

derives from N_0 itself. It is as if N_0 generates *derivative* norms like "ϕ when it is likely on your evidence that you are in C". Moreover, such derivative norms themselves generate further derivative norms. I don't propose any general form for a derivative norm to take. Indeed, there might be derivative norms of various different kinds, each requiring its own story of how it is generated from the primary norm, and what (if any) its normative status is. However, I do want to focus on a particular principle for generating derivative norms. Here is the rough idea: if a primary norm tells one to ϕ in circumstances C, then there is a derivative norm telling one to ϕ when *given one's perspective on the world*, it is as if one is in C. Such a perspective might be captured by one's evidence, by doxastic states based on the evidence, or just by one's doxastic states.[64]

Consider epistemic rules or norms. One might think that the fundamental primary rule in the epistemic case is "Believe *p* just in case *p* is true". Hence, evidential norms might themselves be viewed as derivative. But whether or not this is right (and I am inclined to think that it is not), they generate further derivative norms. I take the fundamental evidentialist norm to be "Proportion your beliefs to your evidence!". I take this norm to entail something along the lines of (E_0):

(E_0) Believe *p* if and only if *p* is (sufficiently) likely on your evidence!

Such a norm generates derivative norms like:

(E_1) Believe *p* if it is likely on your evidence that *p* is (sufficiently) likely on your evidence.[65]
Or: Believe p if you (rationally) believe that p is sufficiently likely on your evidence.

(E_1*) Don't believe *p* if it is likely on your evidence that *p* is not likely on your evidence.
Or: Don't believe p if you (rationally) believe that p is not sufficiently likely on your evidence.

Again, where the primary norm urges subjects to proportion their doxastic states to their evidence, a loose way of characterizing the derivative norms is by saying that they urge

[64] Williamson (*forthcoming*) explores an alternative, though closely related, form for derivative norms to take: "Be disposed to be such as to ϕ in circumstances C".

[65] Or, if one thinks that a subject's perspective is best captured by her beliefs (or degrees of confidence), or perhaps by her *rational* beliefs (or degrees of confidence), rather than by what her evidence supports, then the derivative norms could be formulated in terms of belief. Such an alternative formulation of (E_1) would be "Believe *p* if you (rationally) believe/are confident that *p* is likely on your evidence."

subjects to proportion their doxastic states to their *perspectives* on what their evidence supports – at least when they have such clear perspectives. Or, abstracting away from evidentialism, if the primary norm urges having doxastic states that one ought, in the epistemic sense, to have (whatever the correct theory of epistemic permissions or requirements is), the derivative norms urge having doxastic states that, *from one's perspective*, are the ones that one epistemically ought to have. So, for instance, if a subject has lots of evidence for a false theory of epistemic rationality, and that theory urges her to believe *p*, even though her evidence fails to make *p* likely, then a derivative norm might urge her to believe *p*.

Consider now a subject who has evidence that is radically misled about what it supports. Assume first that *p* is in fact likely on the evidence, but it is also likely that *p* is not likely. Based on her evidence, she becomes confident that her evidence does not support *p*. The intuition to be explained is that there is something wrong with her if she believes *p*, while believing, or being confident, that her evidence does not support *p*. The account I want to propose begins with the observation that in believing *p*, the subject would violate a derivative norm along the lines of (E_1*). Conversely, consider a subject with evidence that does not make *p* likely, but that makes it likely that *p* is likely. Assume that based on this evidence, the subject becomes confident that her evidence supports *p*. If the subject then fails to believe *p*, she violates the derivative norm (E_1). Again, there is no way of abiding by both the primary and derivative norms. Similarly, consider a subject who acts in the seemingly schizophrenic way that an akratic state would seem to rationalize: she bets on *p* at odds 9:1, while betting at odds 1:9 that those odds are irrational. While such a subject respects the norm "Bet on *p* at odds 9:1 if and only if your evidence supports *p* to degree 0.9 or above", she violates the derivative norm "Don't bet on *p* at odds 9:1 if it is likely on your evidence that your evidence doesn't support *p* to degree 0.9 or above". Hence, epistemically akratic subjects violate derivative epistemic norms.

This observation, I am suggesting, is where we should start in order to understand the seeming normativity of the enkratic requirements. Now, one might worry at the outset that the distinction between primary and derivative norms provides insufficient machinery to account for the appeal of the enkratic requirements, for there is an important difference between subjects who are epistemically akratic (and hence, who violate some derivative epistemic norms), and subjects who violate any humdrum derivative norm. Consider a subject who violates a derivative tea-making norm: she doesn't pour the water into the cup, even though she has excellent (albeit misleading) evidence that it is 90°C. However, the subject doesn't recognize that the tea-making norm applies to her, and she is not in any way trying to follow it. Perhaps, for instance, she holds false beliefs about how to make good green tea, believing that the water should be boiling. Then, we feel little inclination to say that she ought to pour the water into the cup, or that it is irrational not to do so. By contrast, doesn't it seem irrational for a subject who believes that her evidence doesn't make *p* likely to believe *p*, whether or not she in any way endorses the evidential norm or recognizes that it applies to her?[66] It is not at all clear to me what the intuitive verdict about such cases is. Above I considered a case in which a subject has evidence pointing to the falsity of the evidential norm itself. Epistemic justification, she thinks, is not about proportioning one's beliefs to the

[66] Thanks to Alex Worsnip for pressing this objection.

evidence. If such a subject believes *p*, while believing (or at least being confident) that her evidence does not support *p*, she does not view her own belief as epistemically unjustified. At the very least, such a background story involving false but evidentially supported beliefs about epistemic justification seems to render the subject who believes *p* while believing that *p* is not likely on her evidence *less* incoherent. The sense that it is irrational in some deep way to believe against one's beliefs about what the evidence supports, irrational in a way not explained by the distinction between primary and derivative norms, might simply result from the fact that it is difficult to imagine subjects who fail to feel the pull of an evidential norm. Perhaps the idea that evidence provides epistemic reasons simply has so much *a priori* appeal that we automatically project onto others some sort of commitment to an evidential norm, whereas we do not automatically project commitments to any humdrum norm.

What I am suggesting is that the right account of the pull of enkratic requirements will fall out the correct story of the seeming normative force of derivative norms. It seems that we can, at the very outset, rule out an account on which derivative norm simply *inherits* the very ought that attaches to the primary norm. Such an account merely restates the original puzzle, for as we have seen, primary and derivative norms can make conflicting recommendations. Moreover, there is something such an account leaves unexplained. In particular, there appears to be something in common to subjects who violate different derivative norms, be these moral, prudential, or epistemic. The subject who sets out to follow the tea-making norm, who has excellent evidence that the water is 90°C, but who fails to pour it into the cup seems to be faulted *in much the same way* as a subject who has excellent evidence that her evidence supports *p*, but who fails to believe, or even be confident, in *p*. It would be desirable to have a somewhat unified account of seeming force of derivative norms. Saying that one subject fails to conform to an epistemic norm, whereas the other fails to conform to a norm of tea-making, does not offer this kind of unification.

I cannot here provide a full story of derivative norms, but will conclude with a promissory note for what form it might take. I will first briefly discuss a view on which the normativity in question is merely apparent, a view that closely resembles Niko Kolodny's (2005) "Transparency Account" of norms of rationality. I then sketch an alternative account, one appealing to dispositions or virtues, on which the normativity in question need not be viewed as merely apparent.

Consider a primary norm urging one to ϕ just in case some condition C obtains, and a derivative norm urging one to ϕ just in case it is from one's perspective as if condition C obtains (whether this should be taken to require having strong evidence that C obtains, believing or being confident that C obtains, or all of these). Assume that a subject has excellent albeit misleading evidence that she is in conditions C, and who believes (or is at least confident of) this. As a result of her misleading evidence, she conforms to the derivative norm, but fails to conform to the primary norm. Now, of course, she believes (or is confident) that the appropriate conditions for ϕ'ing formulated by the *primary* norm obtain. Given the way things seem to her, or given her perspective on the world, the primary norm requires her to ϕ. Hence, there is a sense in which the conditions formulated by the derivative norm are *transparent* to the subject: when they obtain, from the subject's perspective, it is as if the

conditions formulated by the primary norm obtain. Similarly, when a subject's perspective represents it as being the case, for instance, that *p* is very likely on her evidence, it represents the conditions for appropriate belief formulated by the primary evidential norm (E_0) as applying. Perhaps this is enough to explain the apparent normativity of derivative norms: conforming to a derivative norm is what a subject ought to do, assuming that her perspective on the world is accurate.[67]

While I think the kind of transparency account sketched is promising, one might worry whether it can offer a full enough story. For one thing, it is not clear whether we merely want to say that subjects who violate derivative norms are doing something they ought not*, given their perspectives*, be doing, or that akratic subjects are in doxastic states that are epistemically irrational *given their perspectives*. That such subjects fail to follow their own perspectives is something we knew to start out with. And yet, the sense that they are to be faulted might persist.

Another place to look for an account of the force of derivative norms begins with the following observation: normally our only way of reliably conforming to a norm – that is, not only conforming to it on one occasion, but being disposed to conform to it across a range of somewhat normal situations – is by going by our perspectives (our evidence, beliefs, etc.) about whether we are in the relevant circumstances. That is, a disposition to conform to a primary norm is, in paradigm cases, a disposition to conform to derivative norms. For instance, subjects who reliably conform to the tea-making norm base their decisions about when to pour water into a cup on their evidence about the temperature of the water. Indeed, perhaps a standard explanation of how it is that a subject manages to conform to a primary norm is that she does so by conforming to (or following) the derivative norm. In so far as conforming to the primary norm is a good thing, then it seems that having a disposition to conform to it is similarly good – good full-stop, not merely good given a subject's perspective or point of view. A disposition to conform to derivative norms, then, can be thought of as a virtue.

At the very least, an account connecting derivative norms with the virtues of conforming to primary norms is promising in explaining why we are inclined to positively evaluate subjects who conform to derivative norms: in paradigm cases, such subjects manifest a disposition to conform to the relevant primary norms. Nevertheless, one might worry is that such an account cannot provide resources for explaining the seeming *normativity* of derivative norms. In general, the fact that a given act would display a virtue does not give one a reason to do it. For instance, running through a busy street with my eyes shut might provide me with an opportunity to display courage, but this is no reason to run through the street.[68] However, even if we don't automatically have a reason to seek occasions to manifest virtues, we may have a reason to *have* those virtues: if one has a reason to conform to a primary norm, then it

---

[67] While such an account closely resembles Kolodny's Transparency Account of the 'ought' of rationality, it does not draw on many of the controversial assumptions of Kolodny's view. For one thing, I have avoided commitment to a notion of rationality that encompasses the enkratic requirements. I am therefore also not committed to Kolodny's (2005) claim that all norms of rationality can be derived from two narrow-scope core requirements, the first of which requires subjects to have the attitudes they believe they have conclusive reason to have (C+), and the second of which requires subjects *not* to have attitudes they believe they lack sufficient reason to have (C-).

[68] Cf Kolodny (2005: 546-547, 554).

also seems that one has a reason to have a general disposition to conform to the primary norm. Moreover, certain manifestations of a disposition are difficult to detach from the having of the disposition. It is not at all difficult to imagine a courageous person who does not run through busy streets with her eyes shut whenever the opportunity arises. But without some background story, it is more difficult to imagine a courageous person who fails to jump into a lake to save a drowning child. In general, we feel a definite resistance to admitting that one ought to have a certain disposition, while claiming that one ought not to have manifested it on an occasion where the disposition is characteristically manifested.

Here, then, is at least a tempting line of thought: If one ought to conform to a primary norm, then one ought to have a disposition allowing one to reliably conform to it. At least in paradigm cases, this will be a disposition to also conform to the derivative norm. At least in paradigm cases, conforming to a derivative norm is a characteristic manifestation of a disposition to reliably conform to a primary norm. But then, it would seem strange to concede that one ought to have the disposition in question, but not to manifest it by conforming to the derivative norm on particular occasions. Now, because the 'ought' of the derivative norm cannot be the same as the 'ought' of the primary norm, the sense in which one ought to have a disposition to conform to a primary norm must be different from the ought attaching to the primary norm itself. There must be some equivocation going on somewhere in the above line of reasoning. Perhaps the sense in which one ought to have the disposition to conform to the primary norm involves a more instrumental kind of 'ought'. The kind of account just outlined connecting derivative norms with dispositions or virtues to conform to primary norms merely shifts our gaze to the sense in which one ought to have such dispositions. But I think that this may be just the right place to look for an account of the force of derivative norms and, ultimately, for an account of why the enkratic requirements seem appealing.

Such a suggestion is very skeletal, but it already provides resources for saying what is bad, across the board, about epistemically akratic subjects: either such subjects violate primary epistemic norms, failing to proportion their beliefs to their epistemic reasons or evidence (such as subjects who believe in astrology, despite recognizing that there is lots of evidence pointing to its unreliability), or they at least fail to manifest general good dispositions to conform to the primary epistemic norms, thereby lacking a kind of epistemic virtue – indeed, in lacking a disposition to proportion their beliefs to their evidence, they may be positively manifesting an epistemic vice.

Conclusions

The starting point of my investigation was a paradox about epistemic rationality, a paradox that seemed to place what I have called *evidentialism* into direct conflict with the *enkratic requirements* on epistemic rationality. I set out to see whether, and exactly how, the paradox could be solved without rejecting evidentialism. However, I argued at the outset that similar paradoxes can be generated for a wide range of different theories about epistemically rational doxastic states. In particular, essentially the same paradox arises given any theory on which it can sometimes be epistemically rational to hold false beliefs about matters concerning which doxastic states it is rational for one to be in (or at least to place a high degrees of confidence

in falsehoods concerning such matters). In so far as rational belief, or rational high confidence, is not in general factive, why should belief concerning such epistemic matters be an exception? Besides, I argued, we should be wary of the intuitive support that the enkratic requirements seem to enjoy, for there are intuitions that come into direct conflict with them.

The first solution to the paradox that I discussed was denying that evidence can ever be radically misled about what it supports. I argued against such a constraint on possible evidence, but perhaps even more importantly, imposing the constraint would leave a similar paradox unresolved, for it would seem possible to have evidence that is misled about what the correct story of epistemic rationality is in the first place. A second solution was to maintain that though subjects could be *permitted* to be in akratic states, they could never *permissibly be* in such states for defeat-related reasons: the belief, for instance, that it is irrational for one to believe *p* automatically defeats whatever reasons or evidence for *p* one may have had prior to forming this belief. I argued that such a solution fails for fundamentally the same reason as the first: its claim that having a certain kind of body of evidence is metaphysically impossible has counterexamples. Besides, the defeat view lacks resources to rule out the possibility of certain kinds of cases of rational epistemic akrasia. The third solution drew on a distinction between *reasons* and *rationality*, arguing that whereas evidential norms concern the 'ought' of reasons, enkratic norms concern the 'ought' of rationality. Though these two oughts can come into conflict, at least we don't have a conflicting views about epistemic rationality. But accepting the basic distinction between reasons and rationality does not yet suffice to resolve the paradox, for it would also have to be maintained that a subject who perfectly responds to her epistemic reasons can nevertheless be irrational. This strikes me as an unpalatable conclusion with no clear precedents. Besides, I argued, it is not at all clear what the required notion of rationality amounts to.

The remaining solution, I urged, is rejecting the enkratic requirements as genuine requirements on rationality. How, then, should their seeming appeal be explained? I argued that we should turn our gaze to a more general phenomenon, the way in which any norm seems to generate further, derivative norms that sometimes conflict with the original. This raises a general explanatory challenge, not in any way specific to the epistemic case: if derivative norms don't simply inherit the very same 'oughts' attaching to the original ones - and they can't, it seems, for their recommendations sometimes conflict with those of the original norms - what account of their seeming normative force should be give? We should expect our explanation of this phenomenon to generate an account of the appeal of the enkratic requirements, for while an epistemically akratic subject might conform to a primary epistemic norm such as "proportion your doxastic states to your evidence", she will violate various derivative norms (such as "don't believe *p* if you (rationally) believe that *p* is not likely on your evidence"). The methodological reasons for looking for such an account of the appeal of the enkratic requirements are strong, for it would encompass a wide phenomenon that the epistemic case is but one instance of. In the end I provided the beginning of such an account: the problem with akratic subjects is that either they fail to conform to evidential norms, or they fail to manifest general dispositions to conform to such norms, thereby failing to manifest a kind of epistemic virtue.[69]

---

[69] I am grateful to audiences at the University of Oxford, University of Bristol, University of Helsinki, University of St. Andrews, Stanford University, and Rutgers University. Special thanks to Ville Aarnio, Billy

References

Alston, William
(1985) "Concepts of Epistemic Justification", *The Monist*, 68.1: ??.

Bergmann, Michael
(2005) "Defeaters and Higher-Level Requirements", *The Philosophical Quarterly*, Vol 55. 220: 419-436.

Broome, John
1999: "Normative Requirements", *Ratio*,12.4: 398-419.
2013 *Rationality Through Reasoning*, Wiley Blackwell.

Chisholm, Roderick
(1989) *Theory of Knowledge*, third edition, Prentice-Hall.

Christensen, David
(2007) "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals", *Oxford Studies in Epistemology* 2: 3-31.
(2010a) "Higher-order evidence", *Philosophy and Phenomenological Research* 81(1): 185– 215.
(2010b) "Rational Reflection", *Philosophical Perspectives*, Vol 24.1: 121-140.

Coates, Allen
"The Enkratic Requirement", *European Journal of Philosophy*, 21.2: 320-333.

Davidson, Donald
(1985) "Incoherence and Irrationality", *Dialectica* 39.4: 345-354.

Elga, Adam
(2005) "On overrating oneself… and knowing it", *Philosophical Studies* 123: 115-124.
(2013) 'The Puzzle of the Unmarked Clock and the New Rational Reflection Principle', *Philosophical Studies*, 164.1: 127–39.

Feldman, Richard
(1988) "Subjective and Objective Justification in Ethics and Epistemology", *The Monist* 71.3: 405-419.
(2005) "Respecting the evidence", *Philosophical Perspectives* 19.1: 95-119.

Feldman, Richard & Conee, Earl
2004 "Evidentialism", reprinted in E. Conee & R. Feldman *Evidentialism* (Oxford: Clarendon Press) with Afterword, pp. 83-107.

Field, Hartry (2000) "A priorism as an evaluative notion"

Gibbons (2006) "Access Externalism", *Mind*, Vol 115: 457, pp. 19-39.

Dunaway, Branden Fitelson, John Hawthorne, Jennifer Lackey, Clayton Littlejohn, Han van Wietmarschen, and Tim Williamson, and above all to Alex Worsnip for detailed, insightful comments on a previous draft.

Goldman, Alvin
(1986) *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.
(2010) "Epistemic Relativism and Reasonable Disagreement", in R. Feldman and T. Warfield (eds.), *Disagreement*, Oxford University Press, pp.

Greco, Daniel (2014) "A Puzzle About Epistemic Akrasia", *Philosophical Studies* 167.2: 201-219.

Hazlett, Allan (2012) "Higher-Order Epistemic Attitudes and Intellectual Humility", *Episteme* 9.3:205-223.

Horowitz, Sophie (2014) "Epistemic Akrasia", *Nous*, 48.4: 718-744.

Huemer, Michael (2011) "The Puzzle of Metacoherence", *Philosophy and Phenomenological Research*, 82(1): 1-21.

Jackson, Frank
(1991) "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection", *Ethics* 101.3: 461-482

Kolodny, Niko (2005)
"Why be Rational?", *Mind*, Vol 114: 455, pp. 509-563.

Lasonen-Aarnio, Maria
(2010) "Unreasonable knowledge", *Philosophical Perspectives*, Vol. 24:1, pp. 1-21.
(2013) "Disagreement and evidential attenuation", *Nous*, 47.4: 767-794.
(2014) "Higher-order evidence and the limits of defeat", *Philosophy and Phenomenological Research*. 88.2: 314-345.
(2015) "New Rational Reflection and Internalism about Rationality", *Oxford Studies in Epistemology 5*, pp. 145-179. 2015.

Littlejohn, Clayton
(*Forthcoming*) "Stop Making Sense? On a Puzzle About Rationality", *Philosophy and Phenomenological Research*.

Lockhart, Ted (2000) *Moral Uncertainty and Its Consequences*, Oxford: Oxford University Press.

Pollock, John
(1979) "A Plethora of Epistemological Theories", in G. S. Pappas (ed.), *Knowledge and Justification*, Dordrecht, Holland: D. Reidel, pp. 93-113.

Ross, Jacob (2006) "Rejecting Ethical Deflationism", *Ethics* 116: 742-768.

Scanlon, Thomas
(1998) *What we owe to each other*, Cambridge: Belknap Press.

Schroeder, Mark
(2009) "Means-end coherence, stringency, and subjective reason", *Philosophical Studies*, 143.2: 223-248.

Sepielli, Andrew
(2009) "What to Do When You Don't Know What to Do", *Oxford Studies in Metaethics* 4:5-28
(2014) "What to Do When You Don't Know What to Do When You Don't Know What to Do…", *Nous*, 48.3: 521-544.

Shah, N.
(2003) "How Truth Governs Belief" *Philosophical Review* 112 pp. 447-482.
(2006) "A new argument for evidentialism" *The Philosophical Quarterly* 56 pp. 481 – 498.

Smith, Michael
(1994) *The Moral Problem*, Blackwell.

Smithies, Declan
2012 "Moore's Paradox and the Accessibility of Justification", *Philosophy and Phenomenological Research* LXXXV: 2, 273-300.

Titelbaum, Michael G.
(2015) "Rationality's Fixed Point (Or: In Defence of Right Reason)", *Oxford Studies in Epistemology 5*: 253-294.

Wedgwood, Ralph
2002 "The Aim of Belief"
Forthcoming "Objective and Subjective 'Ought'", Nate Charlow & Matthew Chrisman (eds.), *Deontic* Modality, Oxford University Press.

Williamson,Timothy
(2000) *Knowledge and Its Limits*, Oxford: Oxford University Press.
(2014) "Very Improbable Knowing", *Synthese* 79: 971-999.
(*forthcoming*) "Justifications, Excuses, and Sceptical Scenarios", to appear in Julien Dutant & Daniel Dohrn (eds.), *The New Evil Demon*, Oxford: Oxford University Press.

Worsnip, Alex
Forthcoming "Narrow-Scoping for Wide-Scopers", *Synthese*.
Manuscript "The Conflict of Evidence and Coherence"