

Enkrasia or evidentialism? Learning to love mismatch

Forthcoming in *Philosophical Studies*

Maria Lasonen-Aarnio

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 sketch an account that rejects the enkratic requirements, while nevertheless explaining our sense that epistemic akrasia is a distinct kind of epistemic failure. Central to the account is distinguishing between two evaluative perspectives, one having to do with the relevant kind of success (proportioning one's doxastic states to the evidence), the other having to do with manifesting good dispositions. The problem with akratic subjects, I argue, is that they manifest dispositions to fail to correctly respond to a special class of *conclusive* and *conspicuous* reasons.

I A Paradox

Sometimes we lack access to our reasons for action. To take a classic example from Bernard Williams (1979), the fact that a glass I am holding contains petrol is a reason not to drink from it, but I might have good reason to believe that I am holding a glass of gin. At first sight it may look as though no such problem of access arises for *justifying reasons for belief*, as such reasons are already restricted to the subject's perspective on the world, being of a more subjective sort.¹ But alas, matters are not this simple.²

Facts about what we are justified in believing are themselves objective facts, and most epistemologists these days, whether of an internalist or externalist bend, would deny that we have full access to them. Just as there can be misleading evidence about whether a glass 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

¹ For such talk of a subject's "perspective on the world" see, for instance, Alston (1985).

² Against, for instance, Feldman (1988). I sketch a more detailed argument below.

bearing on just what one's evidence supports, or what it is rational to believe given one's evidence.³ 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.

Hence, 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 (I offer more support for the possibility of such situations below). What doxastic state regarding p ought the subject to have in such a situation?⁴ 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, evidentialism is not essential for creating trouble with these requirements, for problems arise given any of a wide range of normative theories in epistemology.

Evidentialism is often characterized as a supervenience claim according to which facts about what doxastic states it is epistemically rational or justified for a subject to be in are fixed by her evidence.⁵ But instead, what I label *evidentialism* is the view that epistemically 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 core claim of evidentialism as I will understand the view, 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 . A simple, somewhat natural way to think is that one is only permitted (and perhaps even required) to have credences that perfectly *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 is intended to be compatible with a view on which credences that are sufficiently *close* to evidential probabilities

³ 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.

⁴ 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).

⁵ See, for instance, Feldman and Conee (2004).

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. Finally, note that though the tracking claim looks to entail, but not be entailed by, the familiar supervenience-thesis, I intend to leave it open whether such an entailment holds. Perhaps, for instance, pragmatic stakes can affect what it is for a full belief to track one's evidence.⁶

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:

(1) p is likely on the evidence, but it is also likely that p is not likely on the evidence.

(2) p is not likely on the evidence, but it is likely that p is likely on the evidence.

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, situation (1). Evidentialism would permit (perhaps even require) 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 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

⁶ 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. Here is a second reason why the entailment might fail. Take 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, though I don't subscribe to such a view, what I say below is even compatible with a radical subjective Bayesian view on which *any* prior probability function is rational.

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. If the beliefs of our akratic subject concerning what she is forbidden or required to believe 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 misleading regarding 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.⁷

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 one's evidence might be misleading when it comes the very question regarding which states track evidential probabilities. 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, that 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 misleading regarding itself entails the falsity of the enkratic principles.

The assumptions that gave rise to the paradox were:

1. Evidentialism
2. Evidence can be radically misleading regarding itself.
3. If 1. and 2., then the enkratic principles are false.

⁷ 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 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, and in numerous places 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 from the concept of rationality. For further defenses of versions of EP-, see also Chisholm (1989: 6), Scanlon (1998: 25), Bergmann (2005), Elga (2005), Coates (2011), and Hazlett (2012).

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 paradox 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 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.⁸ 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 justified or 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 entirety of the above condition. In such a situation, even our new revised theory would

⁸ For a more detailed argument for a similar conclusion, see Lasonen-Aarnio (2014).

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.⁹ What is striking about the paradox is that many have thought that though abiding by so-called *structural requirements of rationality* doesn't entail believing or intending as one's reasons require, at least the converse is true: a subject who believes (and intends) as her reasons require respects structural requirements of rationality.¹⁰ But the paradox seems to show that believing in accordance with one's epistemic reasons sometimes forces one to violate putative instances of structural requirements of rationality, namely, the enkratic requirements.

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 ("Proportion your doxastic states to your evidence!") 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 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.

There is a range of rich, interesting responses to the paradox sketched above that I won't be able to discuss in detail. I won't discuss a view that responds by accepting a pluralism of different notions of epistemically permitted or justified belief: 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), etc.¹² Neither will I discuss a view that accepts both evidential and enkratic norms on belief, viewing situations in which they come to conflict as epistemic dilemmas.¹³ If there is no other way of solving the paradox, we may yet have to resort to such views. However, before seriously scrutinizing the case in favour of the enkratic requirements, I don't think we should be that swift to accept them, thereby rejecting a purely evidentialist view. In the end I show

⁹ 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 being coherent.

¹⁰ E.g. Kolodny (2008).

¹¹ See, for instance, Horowitz (2014) and Titelbaum (2015).

¹² See Sepielli (2014) for a defence of a similar kind of pluralism.

¹³ This is one way of interpreting the kind of view that Christensen (2010a) defends.

how the paradox can be solved without adopting either a plurality of oughts or a plurality of norms that sometimes come into conflict.

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 misleading regarding 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, and say why rejecting it doesn't get to the heart of the problem anyway. 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 misleading regarding 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.

Instead, I argue for a solution that explains the failure of akrasia without recourse to the enkratic requirements. Indeed, the idea of structural coherence that the enkratic requirements are intended to (partly) capture plays no role in the account I propose. What, then, explains our sense that epistemic akrasia is a distinct kind of epistemic failure? I will argue that akratic subjects manifest a failure to function well by exercising basic epistemic competence. Such competence consists in appropriate sensitivity to a special class of *conclusive* and *conspicuous* reasons (or evidence). The kinds of cases discussed above are ones in which subjects can successfully conform to the evidential norm – or, more generally, be in doxastic states that are appropriate given their epistemic reasons – only by displaying a kind of epistemic vice, a failure to exercise competences to appropriately respond to conspicuous reasons.

I want to clear up a couple of potential misunderstandings, at the very outset, concerning the commitments of the strategy I will pursue. 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 think that the first- and higher-order components of a body of evidence can be thus 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. Relatedly, I do not deny that there are no so-called higher-order defeaters.¹⁴ Finally, I am of course not claiming that any state involving epistemic akrasia is acceptable from the perspective of the correct epistemic norms. For instance, a subject who knows that her evidence does not support the truth of astrology, while continuing to believe in astrology nevertheless, fails to proportion her beliefs to her evidence. However, sometimes believing in accordance

¹⁴ For reasons why we shouldn't think that higher-order defeat *always* occurs, see Lasonen-Aarnio (2013, 2014).

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

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 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, and what the notion of rationality being appealed to is. For now I will follow the somewhat standard practice of talking about epistemically justified or permitted belief, on the one hand – which, in an 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 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]¹⁵

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*.¹⁶ The enkratic principles are *wide-scope*

¹⁵ 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.

¹⁶ Note that I am counting believing *p* and believing *q* as distinct types of doxastic states.

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.¹⁷

Recently epistemologists have defended, in particular, some version of the negative requirement EP-.¹⁸ Couldn't they reject the positive requirement?¹⁹ One might worry that such a view is poorly motivated, for numerous arguments and considerations in favor of EP- also favor EP+. Besides, there is a case 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.²⁰ 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.²¹

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).

¹⁷ 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 has enjoyed considerable support (e.g. Kolodny 2008). The K axiom is widely, 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.

¹⁸ See Titelbaum (2014), Horowitz (2015).

¹⁹ 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 .

²⁰ 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.

²¹ 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.

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. 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.

Many have found their favourite enkratic principles so plausible as not to be in need of further defense.²² 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) $\neg p$, though rationality requires me to believe p .
- (2*) $\neg p$, though p is likely on my evidence.
- (3) p , but I am rationally required to suspend judgment on the matter.

But such beliefs and assertions have a Moore-paradoxical feel to them.²³ A way of sharpening the seeming paradoxicality here further is by considering possible cases in which the evidence

²² 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).

²³ See, for instance, Feldman (2005: 108), Bergmann (2005: 424), Hazlett (2012: 211), Huemer (2011), Smithies (2012).

is radically misleading regarding itself in the following way: it makes p likely, while making it likely that it makes $\neg p$ likely.²⁴ Given evidentialism, it looks like a subject with such evidence could believe p , but $\neg p$ 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.²⁵

Moreover, mismatched, akratic doxastic states seem to vindicate mismatched 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 rational. 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 sound strange 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 there are special requirements of rationality forbidding them. Indeed, below I argue that there are combinations of states inviting a very similar charge of irrationality as akratic ones do, even though very few would attempt to explain the seeming badness of the states in question by appeal to some special requirement of rationality. As a result, the argument from Mooreparadoxicality is only as strong as alternative explanations of this seeming badness are weak. 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.²⁶ 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

²⁴ The latter does not follow from the former, but see Horowitz (2014) for a convincing example.

²⁵ See Horowitz (2014).

²⁶ 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 $\neg p$ - and hence, that one's evidence regarding p is misleading.

degrees of confidence. How, for instance, could it be rational for a subject to be *confident* that her evidence regarding a proposition *p* is misleading?²⁷

But, as I argue in more detail below, evidentialists are hard pressed to deny that such states could ever be rational – 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.²⁸ 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?).²⁹ In offering a framework in which the enkratic requirements have no theoretical 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 misleading regarding 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.³⁰ 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 not realize that she is Ama and hence, not realize that *she* believes *p*. Or, she might realize that she is in a given doxastic 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

²⁷ Cf Horowitz (2014: 726).

²⁸ See, for instance, Horowitz (2014), Titelbaum (2015), and Elga (2005).

²⁹ Horowitz (2014) offers a 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.

³⁰ Some proponents of the principles want to restrict them anyway for similar reasons. For instance, Titelbaum (2014) defends EP-, but explicitly restricts the principle to cases that don't involve failures of "state luminosity", thereby excluding precisely the sort of case under consideration in which a subject is not suitably aware of her own mental states.

such revisions, that only strengthens my overall case against them. This worry will be pressed further below.

Given that the enkratic requirements are often defended on grounds of intuitiveness, it is worth flagging 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. 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.³¹ 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.³² 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 .³³ A more concessive response rejects the enkratic requirements in favour of some sort of *rational reflection* principle. I argue against rational reflection principles elsewhere.³⁴ For now I want to focus on the enkratic requirements.

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.

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

³¹ See also Lasonen-Aarnio (2015) for a discussion of the conflict between *rational reflection* principles and the enkratic requirements.

³² 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!

³³ 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.

³⁴ Lasonen-Aarnio (2015).

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. My aim here will be to provide a concise summary of arguments against such a view. Besides, rejecting Premise 2 still leaves open a closely related paradox.

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 failures of the *Positive Introspection Axiom* of modal logic). Pretty much any view of evidence is susceptible to a Williamsonian-style anti-luminosity argument.³⁵ 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* part of my evidence, then it is certain on my evidence that e is not part of my evidence. We need the following: 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 .³⁶ 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 .

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 misleading regarding what

³⁵ See Williamson (2000, Ch 4).

³⁶ q may or may not be identical with p .

it consists in.³⁷ Indeed, these kinds of cases have persuaded numerous recent proponents of the enkratic requirements to restrict their favoured principles.³⁸ 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.³⁹

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.⁴⁰ 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**.⁴¹ 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 misleading regarding 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.

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

³⁷ 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.

³⁸ See, for instance, Horowitz (2014) and Titelbaum (2015).

³⁹ Horowitz (2014) offers a promising diagnosis of what is special about such clock belief-type cases. I don’t think the diagnosis is general enough, though I cannot argue for that here.

⁴⁰ For a more detailed argument for a similar conclusion, see Lasonen-Aarnio (2014).

⁴¹ 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. Then, there is a sense in which the formalism builds in certainty about facts about evidential support. However, I take this to be nothing more than an idealization built into the formalism.

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.⁴² 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.⁴³ 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 of generating yet higher levels of moral uncertainty, for the subject might also be uncertain what the correct way of factoring in uncertainty about the correct first-order moral theories is.⁴⁴

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 misleading regarding 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.

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

⁴² Jackson (1991) discusses and criticizes the first proposal, defending the second – though he formulates these views in terms of moral utility, not happiness.

⁴³ For a discussion of such cases see Ross (2006), Sepielli (2009, 2014), Lockhart (2000).

⁴⁴ Cf. Sepielli (2014).

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 misleading regarding 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, requiring instead a condition entailing that such higher-order beliefs are only justified if true.⁴⁵

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-.

⁴⁵ Littlejohn (2018a) 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.

Evidentialism and the possibility of evidence that is misleading regarding 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 reformulated 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. 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.⁴⁶ Though proponents of this strategy might want to defend only EP- but not EP+, one could try to 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. Would it, for instance, be plausible to claim that a subject can be permitted to hold contradictory beliefs, even if she can never permissibly do so – shouldn't we insist that there is simply no epistemic situation that permits holding contradictory beliefs? Similarly, proponents of enkratic requirements could dig in their heels here: there is no epistemic situation that permits being epistemically akratic.

Second and more importantly, it's not clear 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 how the story would go. Further, the proposal is threatened by counterexamples. Just like a view that

⁴⁶ 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*.

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. 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.⁴⁷ 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.⁴⁸

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.⁴⁹ 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

⁴⁷ Worsnip (2018) also makes this point.

⁴⁸ 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.

⁴⁹ For such a distinction see, for instance, Kolodny (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).

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.⁵⁰

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, many would deny that a subject who believes as her reasons require could flout a structural requirement of rationality such as an enkratic principle.⁵¹ That epistemic requirements could sometimes conflict with, say, the 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.⁵² 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.⁵³ 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.⁵⁴ Believing something, he notes, is taking it to be true. But then, it is difficult to make sense of a state of

⁵⁰ I am grateful to work by Alex Worsnip (2018) for drawing my attention to this option out of the paradox.

⁵¹ See, for instance, Kolodny (2008). 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.

⁵² See, for instance, Broome (2013: 152).

⁵³ 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.

⁵⁴ See Worsnip (2018). Worsnip argues that such states are even more difficult to make sense of than states of practical akrasia.

believing a proposition p , while believing that p isn't at all likely to be true in light of all the available information.

I agree 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 here 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 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. 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.⁵⁵ Drawing a distinction between norms of reason and norms of rationality is far from equivalent to any such claim. Neither is it clear what the notion of rationality as structural coherence is in the first place. However, 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 that of evidential norms. 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'⁵⁶. It also avoids difficult questions about what reason there is to conform to structural requirements of rationality such as the enkratic requirements in the first place. Indeed, even some of the most well-known

⁵⁵ 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.

⁵⁶ Cf. Kolodny (2005: 555-557).

advocates of such requirements are skeptical about whether there is any reason to be structurally coherent in the first place.⁵⁷

VI Rejecting the enkratic requirements (rejecting premise 4)

We are still left with the paradox. Perhaps it is time we re-think commitment to the enkratic requirements themselves. I expressed doubts above about the very notion of rationality *qua* coherence that they are intended to capture. Ultimately, I don't think that such a notion is helpful in structuring the landscape we are in. The sense that there is something bad about being akratic can be explained without invoking special structural requirements that akratic subjects inevitably violate. Before stating my positive account, let me look at two desiderata on the explanatory task at hand.

I argued above that in some situations akratic states are the only ones that correctly reflect one's reasons for belief, for they are the only states that track the evidence. But there are further reasons besides to think that an evidential norm is insufficient on its own as a diagnostic tool of what is wrong with akrasia. Looking into these will uncover important desiderata for the kind of explanatory task at issue. First, compare the following two subjects. Cal believes *p*, even though *p* is unlikely on her evidence. Al has a very similar body of evidence, and he also believes *p*, though *p* is unlikely. However, Al *also* believes that *p* is unlikely on his evidence. Further, this latter belief is perfectly appropriate given his evidence. Even though both Cal and Al believe *p* on insufficient evidence, one might insist that there is an *additional* problem with Al's mental states, given the mismatch they exhibit.⁵⁸ The additional failure cannot be explained by pointing to the fact that Al fails to proportion his beliefs to the evidence, for so does Cal. Related to this point, the failure involved in being akratic appears *distinct* from any old failure to proportion one's beliefs to the evidence, or to take into account one's epistemic reasons for belief. Questions about just what doxastic states are appropriate given a body of evidence can be highly non-trivial to answer given the potential complexity of evidence. The charge of a blatant kind of irrationality comes less naturally when considering a subject who fails to proportion her beliefs to a complex body of evidence than it does when considering a subject who is akratic.⁵⁹

I will now outline an evaluative perspective focused on *good dispositions* which, I will argue, explains the seeming badness of being akratic, and that satisfies the above two desiderata. Anything from penalty kicks to political campaigns can be evaluated from this perspective, though here my focus will be on evaluating doxastic states. I think this is a perspective we often occupy; in effect, I think it underlies a very wide range of verdicts circulating in the epistemology literature. The perspective in question evaluates doxastic states by evaluating the dispositions manifested in forming and retaining them. The goodness of a disposition, in turn, is always relative to a relevant success: a given disposition is good relative to a success if it tends to manifest itself as doxastic states (or actions) that are successful in this sense.

⁵⁷ E.g. Broome (2013).

⁵⁸ Cf. Kolodny (2007: 237).

⁵⁹ Cf. Way's (2018) discussion of structural rationality.

I will first outline the dispositional evaluative perspective. I then apply it to the puzzle we have encountered, arguing for the following diagnosis of what is wrong with being akratic: given assumptions about feasible ways for humans to be, in failing to respond to a content of her own belief, the akratic subject manifest a bad disposition, a disposition to be unresponsive to a special certain class of reasons or evidence that is *conclusive*, and, in particular, *conspicuous* in what it points to. That is so even if she is in fact perfectly responsive to her reasons or evidence. Sometimes succeeding to be in states called for by the evidential norm – or succeeding to be in states that are appropriate given the reasons for belief one actually has – comes at the cost of manifesting bad dispositions. More generally, success can sometimes even require manifesting bad dispositions, dispositions that lead to failure across a relevant range of counterfactual cases.

The dispositional perspective evaluates the dispositions that a doxastic response such as coming to believe a proposition p is a manifestation of. The *first* step, then, is to determine what dispositions are manifested when one forms or retains a doxastic state, or when one performs an action or series of actions. For instance, if I shoot a penalty kick, what disposition is the kick a manifestation of? Or, if I come to believe p , what disposition is my coming to believe p a manifestation of?⁶⁰ Most pertinent to the present puzzle is the following question: if I retain a belief in p despite believing that my belief in p is irrational, what disposition does this manifest? Am I manifesting a problematic kind of insensitivity to the contents of my beliefs, an insensitivity that tends to lead me astray on other relevant occasions?

The *second* step is determining how good the dispositions being manifested are – in particular, how good they are relative to a relevant success, such as belief proportioned to one's evidence, or knowledge. The goodness of a disposition depends on what tends to happen when it manifests itself and hence, on what happens across a relevant range of counterfactual cases. A good disposition tends to manifest itself as a successful belief or action, or does so at least in a sufficiently high proportion of cases. Take, for instance, the dispositions a soccer player manifests when shooting a penalty kick: she deploys a given technique and strategy, driving a very fast and powerful kick right down the middle of the goal. Now look at counterfactual cases in which these dispositions are manifested. Not all cases should be weighted equally: abnormal cases, or cases that are very different from the actual one in relevant ways, should be assigned no, or very little, weight.⁶¹ (In addition, we might, for instance, be interested in how good the dispositions in question are when playing against a particular goalie. In that case, we only consider cases involving that very goalie. I think dispositional evaluations are highly sensitive to various contextual factors.) Irrespective of whether our subject goals, the dispositions she

⁶⁰ Sometimes people do things that are out of character. Similarly, sometimes we manage our doxastic states in ways that don't manifest general dispositions. However, I will assume that whenever a subject manages her doxastic states in a certain way, she is at least manifesting some more *local* disposition. Assume, for instance, that I retain a belief despite receiving testimony to the contrary from a reliable source. I might not be generally disposed to ignore relevant testimony. Still, I can be manifesting a disposition to ignore testimony from a certain source, or to be obstinate when in a certain emotional state, etc.

⁶¹ Similarly, consider a disposition to form beliefs on the basis of how our experiences represent the world as being. When evaluating the goodness of such a disposition, we normally assign little if no weight to cases involving Cartesian deceivers.

manifests are good assuming that they manifest themselves as successful penalty shots in a sufficiently high range of relevantly similar counterfactual cases.⁶²

Let me now apply this dispositional evaluative perspective to our akratic subject with mismatched evidence. Let the relevant success be that of having beliefs that are appropriately responsive to one's (epistemic) reasons. (Below I will mostly deploy the ideology of reasons instead of that evidence, in order to speak to the more general version of the paradox outlined at the outset.) In what follows, I will be making three assumptions about normative reasons in general, and epistemic reasons in particular. First, not all reasons are mental states: at least some reasons are propositional.⁶³ Though I cannot argue for the view here, it has plenty of support, and many epistemologists are attracted to it.⁶⁴ Moreover, a non-mentalist view of normative reasons is the standard view in meta-ethics. Second, p cannot be a normative reason to ϕ if p is false: falsehoods may serve as merely apparent reasons, but not as genuine reasons.⁶⁵ Third, I shall assume that there is a distinction between normative reasons and what I will refer to as *possessed normative reasons*, or reasons a subject *has*.⁶⁶ For a subject to possess p as a normative reason, she must bear some appropriate epistemic relation to p . It might, for instance, be that of knowledge, though nothing I say rests on this. What I will assume is that the relation entails believing p , or at least being reasonably confident in p . For instance, that Su is in Uppsala might be a reason to believe that she is in Sweden, but unless I bear some epistemic relation to the proposition that Su is in Uppsala (minimally, unless I believe it), it is not a reason *I have* to believe that she is in Sweden.⁶⁷ I find it plausible that p is a possessed normative reason just in case p is part of one's evidence, and that it is part of one's evidence just in case one knows p . However, below I won't assume this.

There is great diversity among reasons to believe a single proposition p ; similarly for reasons to not believe p . I now want to focus on a special class of reasons (or evidence). Consider the reason *I am required, given my evidence, to believe p* . On my view, this is a reason to believe p . If anything does, it certainly counts in favor of believing p !⁶⁸ Moreover, it is not

⁶² Even given a weighting, we will need to balance cases of success against cases of failure. For instance, in determining how good certain dispositions are given the success of knowing, we will need to balance cases in which they manifest themselves as knowledge against cases in which they manifest themselves as false belief – and possibly also cases in which they manifest themselves as suspension of judgment. This might involve deploying something like a *scoring rule*.

⁶³ In so far as there is a distinction to be drawn between facts and true propositions, some non-mentalists might think that it is facts, not propositions, that are candidates for normative reasons. In what follows, I will speak as though *propositions* can be epistemic reasons for belief, though the structure of the account offered does not essentially rely on taking propositions, as opposed to facts, to be such reasons.

⁶⁴ See Sylvan (2016) for a good discussion.

⁶⁵ For an overview of the debate, and a defense of the idea that *all* normative reasons are facts or true propositions, see Littlejohn (2018b).

⁶⁶ For the terminology, see Sylvan (2016).

⁶⁷ To possess p as an epistemic reason to believe q , some further condition might be needed, such as sensitivity to the relation between p and q (cf. Sylvan 2016: 368).

⁶⁸ Some might object: *I am required, given my evidence, to believe p* is not itself a reason to believe p . The objection might draw from arguments made in defense of so-called buck-passing accounts of goodness, or of moral rightness – on the latter accounts, for instance, the moral rightness of an action is not a reason to perform an action. I am unconvinced by these arguments; for an excellent criticism of buck-passing accounts, see Johnson King (*forthcoming*).

just any old reason, for it is *conclusive*: there couldn't be other epistemic reasons that outweigh it in the overall balance of reasons. I may acquire new evidence bearing on *p* that requires me to change my beliefs, but were I to acquire such new evidence, *I am required, given my evidence, to believe p* could no longer be a genuine reason – recall the assumption that falsehoods cannot serve as reasons for belief. On the other side, consider *that p is unlikely on my evidence*. At least assuming some form of evidentialism, this is a conclusive reason not to believe *p*. Of course, similar points could be made without any evidentialist commitments. Consider the reason *I am required, given my epistemic reasons, to believe p*, or simply *I am rationally required to believe p*. These are conclusive reasons to believe *p*. Similarly, *I am forbidden, given my epistemic reasons, to believe p*, or simply *it is irrational for me to believe p*. These are conclusive reasons not to believe *p*.

The kinds of reasons pointed to are not only conclusive, but they are special in another respect as well, for there is a sense in which they wear their force on their sleeves. It is sometimes hard to figure out what one's evidence points to, but it is not that hard to figure out that *that p is unlikely on the evidence* doesn't point to *p*'s truth. The force of such reasons is *conspicuous* in a way that the force of many of our reasons is not. At least without some elaborate background story, failing to appropriately take such conspicuous reasons into account appears blatantly vicious: in the absence of additional information, we expect epistemic agents to appropriately respond to such reasons. Assume that Al cannot quite figure out what his evidence points to. Understandably so, given its complexity. An omniscient being then tells him that *p* is highly unlikely on his evidence. Al believes what he is told. If he then (still) believes *p*, it is very hard to make sense of him: it will certainly no longer do to point to the complex and opaque nature of his evidence. Al's failure at this point looks distinct and more blatant than his initial failure, for he now has a conspicuous reason to not believe *p*. Dispositions to be appropriately sensitive to the kinds of conclusive and conspicuous reasons just discussed will play a central role in my account of the failure of akrasia.

My general suggestion will be that the seeming badness of akrasia be explained in terms of the badness of failing to appropriately respond to conspicuous reasons – or, more precisely, the badness of dispositions to fail to appropriately respond to such reasons. Various accounts of conspicuousness are compatible with my proposed explanation. An account of conspicuousness might appeal to special epistemic features of these reasons: perhaps we have a special kind of epistemic access to what these reasons support, where the access might be *a priori*, or conceptual/analytic, flowing from an understanding of the relevant concepts, etc. Or, an account of conspicuousness might appeal to our expectations as evaluators: absent background information, conspicuous reasons are reasons that we expect any minimally competent subject to be able to appropriately respond to.⁶⁹ Relatedly, it may be that the more conspicuous a reason is, the more of an explanation we need in order to rationalize, or make sense of, a subject who fails to appropriately respond to the reason in question.

Note that it is no part of my view that in general, having doxastic states that are appropriate given one's reasons or evidence requires holding higher-order beliefs about what it

⁶⁹ Cf. Setiya's (2004) notion of a legitimate expectation. Setiya discusses legitimate expectations that an agent will conform to requirements of practical rationality. On my account, of course, such requirements play little role. Rather, what is at issue is legitimate expectations that one respond, and be disposed to respond, to conspicuous reasons.

is rational (irrational) to believe, or about what is made likely (unlikely) by one's evidence.⁷⁰ The vast majority of time we respond directly to our reasons or evidence, without forming such higher-order beliefs. This is perfectly compatible with the existence of such conspicuous reasons, and the thought that failing to appropriately respond to them, *when one has them*, appears blatantly bad.

Let me now put the different parts of my account together. Consider the akratic subject with mismatched evidence: her evidence makes *p* likely, but also makes it likely that it does not make *p* likely. Assume that the subject believes *p*, while believing that it is irrational for her to believe *p* (and, perhaps, that her evidence doesn't support *p*). Her evidence is in fact misleading regarding such higher-order matters. Hence, propositions like *my evidence doesn't support p*, or *it is irrational for me to believe p* are not genuine reasons our subject has, for they are false. So there is a sense in which the subject does well by ignoring these beliefs. But now consider the disposition she manifests by being thus unresponsive to the contents of her beliefs. She manifests at least a local disposition to believe *p*, despite believing that it is irrational for her to believe *p*. How does this disposition fare across counterfactual cases in which it manifests itself? Well, it doesn't fare very well across those counterfactual cases in which the contents of the subject's beliefs in fact constitute genuine reasons. Across a wide range of cases in which our subject believes propositions like *my evidence doesn't support p*, or *it is irrational for me to believe p* on the basis of good evidence, the contents of these beliefs constitute genuine reasons she has. In such cases a disposition to disregard the contents of beliefs such as *it is irrational for me to believe p* manifests itself as unresponsiveness to genuine reasons one has. And what seems particularly bad about such unresponsiveness is that that the reasons in question are conclusive and conspicuous: it appears particularly vicious to manifest dispositions that often manifest themselves as unresponsiveness to conspicuous reasons.

Hence, a disposition to believe *p*, while believing *it is irrational for me to believe p*, or *my evidence does not support p*, is problematic. Why? It fails to discriminate between cases where *it is irrational for me to believe p* is merely an apparent reason, and ones in which it is a genuine reason. In cases in which one *has* the conspicuous reason in question, a disposition to be unresponsive to the content of the relevant belief manifests itself as unresponsiveness to a genuine conclusive and conspicuous reason one has. From the akratic subject's actual way of handling her doxastic states, we can read off a disposition to be unresponsive to a class of conclusive and conspicuous reasons – in fact, the subject manifests a disposition to be thus unresponsive – and such unresponsiveness looks blatantly bad.

In general, the only way in which a subject can be disposed to be responsive to genuine reasons she has is by being responsive to the contents of her beliefs. Couldn't our akratic subject be manifesting a disposition to be unresponsive to the contents of beliefs like *my evidence does not support p* only when those beliefs are false? To answer this question, we need to look at feasible ways for human-like subject to be. Just as humans cannot be disposed to only believe truths, humans cannot be disposed to respond in different ways depending on whether their

⁷⁰ By contrast, Kolodny (2008) argues that a disposition to believe as reason requires has two constitutive parts: a disposition to form true beliefs about what reason require, and a disposition to form or adjust doxastic attitudes in response to one's beliefs about what reason requires – that is, a disposition to conform to something very much like the Enkratic Requirements (or what Kolodny 2008: 457 calls *Believed Reasons*).

evidence is in fact misleading.⁷¹ For instance, if I am disposed to trust a reliable source (when I lack reason to distrust it on a particular occasions), I am disposed to trust it *even when* what it says is false (and I lack reason to distrust it on this occasion). A disposition to only believe p when one has ample evidence for p and p is true, but not when one has ample evidence for p and p is false, would have to discriminate between cases depending only on whether or not p is true. Our akratic subject has misleading evidence for propositions like *my evidence does not support p* and *it is irrational for me to believe p* , and is at least fairly confident in them. But she cannot be disposed to be unresponsive to her belief/confidence in these propositions *only when* the evidence for them is misleading. To have such a disposition would require to be able to track the truth in some direct way not mediated by one's evidence. And this is not something human subjects can be disposed to do.

More generally, assume that p is a reason to ϕ . For instance, *that someone is drowning* is a reason for a lifeguard to jump into the water. Assume that a lifeguard experiences a very realistic illusion as of someone drowning: though she comes to believe that someone is drowning, she fails to jump into the water. Because her belief is false, she doesn't have a genuine reason to jump in. But now consider the dispositions she is manifesting. She cannot be disposed to *only* take action when she believes that someone is drowning, and her evidence is not misleading. In failing to jump in, she manifests a disposition that is indiscriminate between cases in which someone really is drowning, and cases in which she merely falsely believes this based on misleading evidence. And this is a very bad disposition for a lifeguard to have: the lifeguard is not disposed to save people who are drowning. Across a significant range of counterfactual cases, the disposition leads to failing to take action to save people who are drowning. Note also that it's not a stretch to use normative language when evaluating the lifeguard who (falsely) believes someone to be drowning, but makes no effort to save them: 'You should've jumped in – you thought someone was drowning!'. The fact that the lifeguard failed to jump in is evidence that she lacks a competence that is crucial for performing her job well.

I have argued that our akratic subject manifests a problematic disposition: a disposition to be unresponsive to the contents of certain beliefs, a disposition that fails to discriminate between cases in which she has a merely apparent reason due to misleading evidence, and cases in which she has a genuine conclusive and conspicuous reason. The account I have proposed locates akrasia in a larger class of failures involving the manifesting of problematic dispositions. The sense that akrasia is a distinct kind of failure is explained by the special nature of the kinds of reasons discussed above – in particular, their conspicuousness. The distinctness of the failure to appreciate the force of such reasons is one of degree, not kind. But this is as it should be. Indeed, my view gains support from the realization that the charge of blatant irrationality itself comes in degrees. Al believes that a reliable newspaper reported that p , and that there is no reason to distrust the report. Nevertheless, he believes not- p . This is somewhat baffling: isn't Al just irrational? He resembles an akratic subject: perhaps his irrationality is slightly less blatant, but it is still distinct from the failure of a subject whose evidence is complex and

⁷¹ Of course, I can be disposed not to trust bad sources. For instance, I can be disposed not to trust Breitbart News, but this is not a case in which I pick up on the fact that my evidence is misleading: given all my background knowledge about this news source, even if Breitbart News reports that p , I don't have evidence for p . Hence, my evidence is not misleading regarding p .

difficult to decipher. It's just not that difficult to figure out that *a reliable newspaper reported that p, and there is no reason to distrust the report* is a reason to believe, or at least be confident in, *p*. The problem with AI is that he fails to respond to a fairly conspicuous reason.

At this point it is worth distinguishing my account from others floating around in the literature. Various authors have tried to explain the seeming normativity of structural requirements of rationality by using the language of dispositions. A common theme has been the idea that being structurally rational is a *means* of doing what one has reason to do.⁷² Kolodny expresses the idea, which he attributes in some form to both John Broome and Michael Bratman, as follows:

...by being *disposed* to satisfy requirements of formal coherence *over the long run*, one takes means to believing or choosing what reason requires *over the long run*. Surely, if one is disposed to avoid contradictions, or to avoid means-ends incoherence, then one tends to believe and intend more of what reason requires. (Kolodny 2008: 442)

Similarly, one might suggest that having a disposition to be enkratic is good, for over the long run it leads to conforming to the evidential norm. But focusing on such contingent means has the danger of distancing us from the sense that akrasia is a distinctively epistemic kind of failing. After all, getting enough sleep and doing yoga might be means to believing as one ought.⁷³

By contrast, I have not argued that a disposition to conform to the enkratic requirements is merely an instrumental, contingent means to believing what one has reason to believe.⁷⁴ The badness of the dispositions manifested by an akratic subject who is insensitive to the contents of her beliefs is not merely a matter of failing to pursue contingent means to being in doxastic states she has reason to be in. On the view I have proposed, the relevant kind of goodness of a disposition is not a matter of what the disposition tends to promote in a merely instrumental kind of way, or of what its downstream consequences tend to be. It is more direct: good dispositions tend to *manifest themselves as* successful beliefs or actions. For instance, the dispositions manifested by an expert soccer player's penalty kick manifest themselves as a successful penalty kick – that is, as a goal – across sufficiently many counterfactual cases. The problem with a disposition to fail to be responsive to beliefs like *it is irrational for me to believe p* is that this disposition manifests itself across a wide range of counterfactual cases as retaining a belief that is unresponsive to a conclusive and conspicuous reason.⁷⁵

⁷² Kolodny (2008: 442). Bratman (1987: 35) and Broome (2005) and express thoughts along similar lines.

⁷³ Cf. Horowitz (*forthcoming*).

⁷⁴ Note also that unlike Kolodny (2008), I am not even concerned with a disposition to be enkratic “as such”, construed as a disposition to randomly resolve states of akrasia by moving to some enkratic state – dispositions to respond to conclusive and conspicuous reasons don't involve such dispositions. Kolodny is concerned with accommodating the intuition that a subject is doing something right when she resolves a structurally incoherent states, *no matter how she does this*. I cannot here fully discuss the matter, but I have trouble feeling the force of the intuition Kolodny appeals to. Moreover, I think it is easy to confuse judgments about transitions with judgments about end-states: even if there is something superior about any enkratic state as compared to an akratic one, this does not entail that the transition itself leading to the enkratic state is to be evaluated positively, no matter what course it takes.

⁷⁵ Hence, my account bears a resemblance to views on which various forms of structural incoherence are a matter of poor functioning of our abilities to respond to reasons, such as Raz's (2005) view of so-called instrumental irrationality. While Raz struggles to explain what seems distinctive about instrumental

I now want to address two objections which will help me further elucidate the account.

Objection 1: Would it be any *better* for the akratic subject to instead manifest a disposition to be responsive to the contents of beliefs such as *it is irrational for me to believe p*? Giving up her belief in *p* in response to this higher-order belief would lead her to respond to a merely apparent reason, and to fail to be responsive to her reasons or evidence in favour of *p*. How can we criticize someone for failing to be responsive to the contents of her beliefs, if being thus responsive would in fact involve epistemic failure ?

In response, there is an important point I want to emphasize. First, some dispositions manifested by a subject can be problematic *even if* any alternative, feasible way of managing her beliefs (or acting) would likewise be problematic. Indeed, I think that some epistemic situations are simply unfortunate in that there is no humanly feasible way of both succeeding and *not* manifesting some bad dispositions; there might even be cases in which any humanly feasible dispositions one could manifest are problematic. In fact, I think that our akratic subject with mismatched evidence is in a kind of dilemmatic situation: either she will fail to respond to her first-order reasons for believing *p* (and hence, also manifest at least a local disposition to fail to thus respond), or she will manifest a problematic disposition that, across a wide range of counterfactual cases, manifests itself as a failure to appropriately respond to conclusive and conspicuous reasons. Success comes at the cost of a kind of manifesting a kind of vice. But avoiding that vice would come at the cost of failure to respond to actual reasons one has.

As many authors have pointed out, the kinds of evaluations that go together with imputing some form of structural irrationality are *local*: we might, for instance, criticise a subject for failing to pursue means to her ends, even while recognizing that overall this isn't such a bad thing, given that she has very bad ends. On the account I have sketched, the badness of being akratic involves manifesting *some* problematic dispositions, dispositions that tend to manifest themselves as insensitivity to conclusive and conspicuous reasons. *Even if*, all things considered, an akratic subject in fact does better, by believing *p* while falsely believing *my evidence does not support p*, she is still manifesting a problematic disposition, for in many cases a disposition to retain thus mismatched beliefs manifests itself as unresponsiveness to real reasons.

Objection 2: My account relied on assumptions about feasible ways of humans to be: a human-like subject cannot be disposed to be responsive to her belief in a proposition *p* in a way that discriminates between cases in which *p* is likely on her evidence but false, on the one hand, and cases in which *p* is likely on her evidence but true. However, not all cases of akrasia involve beliefs based on good evidence. Consider the following example. Al has excellent evidence to believe that he will pass his final exam, and he in fact believes this. He then goes to see a fortune teller who tells him that he holds an irrational belief concerning some important matter. Based on this, Al comes to believe that his belief that he will pass the exam is irrational. So Al is akratic: he believes that he will pass the exam, while believing that this belief is irrational. However, being a somewhat sensible and responsible subject, Al has a cognitive mechanism

irrationality, ultimately conceding that there is nothing distinctive about it, I have suggested that what is distinctive about akrasia, as opposed to any old case in which a subject's reasons-sensitive capacities malfunction, is that in the akrasia case the relevant dispositions involve sensitivity to conclusive and conspicuous reasons.

that isolates the belief formed as a result of seeing the fortune teller from the rest of his beliefs, as well as from his actions. The result of the mechanism is that Al might still be disposed to treat the proposition *it is irrational for me to believe that I will pass the exam* as a conclusive and conspicuous reason to not believe that he will pass the exam in counterfactual cases in which it really *is* a genuine reason he has, cases in which he has good evidence to believe it, for in those cases the protective cognitive mechanism would not kick in. Al is akratic, and we should be able to explain the seeming badness of this, but it is not clear whether he is manifesting bad dispositions, dispositions that lead him to overlook conclusive and conspicuous reasons across a wide range of counterfactual cases.

However, I don't think we should take it for granted that Al genuinely believes the proposition *it is irrational for me to believe that I will pass the exam* in the example described. In particular, it is contentious whether Al holds a genuine belief, given a wide range of different views on which it is constitutive of belief that it has a certain kind of functional profile. Due to the operations of the cognitive mechanism described, Al's putative belief is assumed to be cognitively isolated from other beliefs, and from actions. Here is what I find to be a very appealing, albeit rough, characterization, of the functional profile of belief: believing *p* looks, from the inside, like possessing the genuine reason that *p*. That is, belief plays the same local role as the epistemic state required for possessing a genuine normative reason – for instance, if possessing *p* as a reason requires knowing *p*, then belief plays the same local role as the state of knowing.⁷⁶ A subject treats the contents of her own beliefs in the same way as she treats genuine reasons. If Al isn't disposed to treat the content *it is irrational for me to believe that I will pass the exam* at all in ways that he would treat this content if it constituted a genuine reason, then it is contentious whether he genuinely believes it.

One could push back here. Treating a content in a certain way – and, relatedly, a mental state playing a certain functional role – is something that comes in degrees. Perhaps we could describe the example so that the relevant state plays *just enough* of the functional role of belief to count as a genuine belief. Or, perhaps it could be argued that it does have the dispositional profile of a genuine belief, it's just that these relevant dispositions are being masked by the cognitive mechanism. But at this point it is worth taking a step back to reflect on what is so puzzling about epistemic akrasia in the first place, and of why it is difficult to make sense of akratic subjects. The less of the functional profile of belief a representation plays, the less it seems like an instance of seriously and in earnest representing the world to be a certain way. At least to me, Al's akrasia looks less bad precisely because his belief in *it is irrational for me to believe that I will pass the exam* is so isolated, and has very few of the powers of a genuine belief.

Let me warp up. We often deploy an evaluative perspective concerned not with success, but with the manifestation of good dispositions, dispositions that tend to manifest as successful doxastic states or actions. Su is the star striker on her soccer team. She gets a stellar pass, but seems to be immersed in her thoughts, oblivious to the ball hurtling in her direction. Still in her trance, the ball happens to bounce off her foot, with the goalkeeper of the opposite team so baffled at what is going on that it rolls straight into the goal. Su's team goaled, surely a good result for them! But what was going on with Su? It's not a stretch to say "You should have

⁷⁶ Cf Williamson (2017).

focused on the game – what were you thinking?” The fortuitous goal doesn’t mitigate such negative evaluations – and note that it is very natural to use the normative vocabulary of should and oughts in such situations. Similarly, even if an akratic subject manages to be in doxastic states that track her evidence, we are baffled, as in failing to respond to the contents of her beliefs – contents that, if true, constitute conclusive and conspicuous reasons to be or fail to be in certain doxastic states – she manifests bad dispositions, dispositions that involve ignoring her reasons of evidence across a wide range of cases. We have here a case in which the only way of conforming to an evidential norm, or of correctly responding to one’s reasons, is by doing something that is vicious from the perspective of conformity to such a norm across a range of other cases. One-off success is achieved by manifesting a bad disposition. That’s why akrasia looks bad across the board.

The primary task I set out for myself was explanatory: to explain our sense that there is something wrong with an akratic subject, even in situations in which her beliefs succeed by being proportioned to the evidence. To provide such an explanation, it is enough to argue that akratic subjects manifest bad dispositions, and that we easily deploy an evaluative perspective focused on the dispositions manifested. The explanation is compatible with the conclusion that deploying the language of oughts in this connection encodes a fundamental error. But it is also compatible with the view that, along with two evaluative perspectives, we have to accept two distinct though intertwined oughts, the ought of success and that of dispositions to succeed. I think this is the right way to go. But arguing for such a claim goes beyond my present aims: pointing to our inclination to respond to instances of manifesting problematic dispositions by criticizing subjects – sometimes even by using normative language – is enough for the explanatory task at hand.

My approach has been somewhat broad-brushed. However, I hope to have shown that an evaluative perspective focused on dispositions can provide an account of what seems bad about akrasia, even if the akratic subject believe as her epistemic reasons require. Cases involving the kind of ‘mismatched’ evidence described above are puzzling, for in such cases two evaluative perspectives – one having to do with one-off success, and the other having to do with manifesting good dispositions – pull us in different directions. Success comes at the cost of manifesting bad dispositions.

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*. I set out to investigate whether 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 normative theories in epistemology. In particular, essentially the same paradox arises given any theory on which one can sometimes be epistemically permitted to hold false beliefs about matters concerning which doxastic states one is permitted 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? What is striking about the paradox is that many have thought that while abiding by so-called *structural requirements of rationality* doesn’t entail believing

or intending as one's reasons require, at least the converse is true: a subject who believes (and intends) as her reasons require respects structural requirements of rationality.⁷⁷ But given plausible assumptions about epistemic reasons for belief, sometimes believing as one's reasons require can force one to be akratic and hence, structurally irrational.

Instead of appealing to special structural requirements of rationality that akratic subjects flout, I have located our sense that there is something wrong with akrasia, even in cases in which akratic states are called for by the evidence, in that fact that akratic subjects manifest bad dispositions. In particular, they manifest dispositions to fail to be responsive to a special class of conclusive and conspicuous reasons. While akrasia is part of a larger class of failures, the special nature of these reasons can explain our sense that akrasia is a *distinct* kind of failure, and that it can be *additional* to a failure to exercise sensitivity to any old reasons or evidence. My account gains further support from the observation that the charge of blatant irrationality we are prone to make when encountering akratic subjects itself comes in degrees, the degrees corresponding to how conspicuous the relevant reasons are.

Investigating whether the kind of account sketched can be extended to explain away the appeal of other structural requirements of rationality is a topic I wish to take up in further work.⁷⁸

References

Alston, William

(1985) "Concepts of Epistemic Justification", *The Monist*, 68.1: 57-89.

Alvarez, Maria

(2010) *Kinds of reasons: An essay in the philosophy of action*, Oxford: Oxford University Press.

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.

(2005) "Does Rationality Give Us Reasons?", *Philosophical Issues* 15: 321-337.

(2013) *Rationality Through Reasoning*, Wiley Blackwell.

⁷⁷ E.g. Kolodny (2008).

⁷⁸ This article took a long time to come out, the initial draft having been written in 2014, and work being interrupted by two babies. I am grateful to audiences at the University of Oxford, University of Bristol, University of Helsinki, University of St. Andrews, Stanford University, Rutgers University, Princeton University, and an APA session. I am grateful to many brilliant philosophers for feedback on earlier drafts. Special thanks to Ville Aarnio, Jessica Brown, Sarah Buss, Stewart Cohen, Billy Dunaway, Branden Fitelson, John Hawthorne, Sophie Horowitz, Nick Hughes, Yoaav Isaacs, Jennifer Lackey, Clayton Littlejohn, Errol Lord, Ofra Magidor, Giacomo Melis, Kieran Setiya, Teruji Thomas, Jonathan Way, Daniel Whiting, Han van Wietmarschen, and Tim Williamson – and special thanks to Alex Worsnip for detailed and insightful comments on a previous draft.

- Bratman, Michael
(1987) *Intention, Plans and Practical Reason*, Harvard University Press.
- 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”, in *New Essays on the A Priori*, Boghossian and Peacocke (eds.), Clarendon Press, Oxford, pp. 117-148
- Gibbons
(2006) “Access Externalism”, *Mind*, Vol 115: 457, pp. 19-39.
- 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.

(Forthcoming) "Predictably Misleading Evidence", in *Higher-Order Evidence: New Essays*, Mattias Skipper and Asbjørn Steglich-Peterson (eds.), Oxford University Press.

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

Johnson King, Zoe A.

(forthcoming) "We Can Have Our Buck and Pass It, Too", *Oxford Studies in Metaethics*, vol 14.

Kolodny, Niko

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

(2007) "How Does Coherence Matter?", *Proceedings of the Aristotelian Society*, 107.1: 229-263.

(2008) "Why be Disposed to be Coherent?", *Ethics* 118.3: 437-63.

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

(2018a) "Stop Making Sense? On a Puzzle About Rationality", *Philosophy and Phenomenological Research* 96.2: 255-513.

(2018b) "Reasons and Theoretical Rationality", in Daniel Star (ed.), *Oxford Handbook of Reasons and Normativity*, Oxford: Oxford University Press.

Lockhart, Ted

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

Marcus, Ruth B.

(1990) "Some revisionary proposals about belief and believing", *Philosophy and Phenomenological Research*, 50: 132-153.

Parfit, Derek

(2001) "Rationality and reasons", In D. Egonsson, J. Joeseffson, B. Petersson, & T. Rønnow-Rasmussen (Eds.), *Exploring practical philosophy: From action to values* (pp. 19-39). Farnham: Ashgate.

Pollock, John

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

- Raz, Joseph
(2005) "The Myth of Instrumental Rationality", *Journal of Ethics and Social Philosophy* 1/1.
- 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.
- Setiya, Kieran
(2004) "Against Internalism", *Nous* 38.2: 266–298.
- 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.
- Sylvan, Kurt
(2015) "What Apparent Reasons Appear to Be", *Philosophical Studies* 172:587–606.
(2016) "Epistemic Reasons I: Normativity", *Philosophy Compass* 11/7 (2016): 364–376
- Titelbaum, Michael G.
(2015) "Rationality's Fixed Point (Or: In Defence of Right Reason)", *Oxford Studies in Epistemology* 5: 253-294.
- Way, Jonathan
(2018) "Reasons and Rationality", in Daniel Star (ed), *Oxford Handbook of Reasons and Normativity*, pp. 485-503
- Wedgwood, Ralph
(2002) "The Aim of Belief", *Philosophical Perspectives* 16: 267-97.

Williams, Bernard

(1979) "Internal and External Reasons", in Ross Harrison (ed.), *Rational Action*, Cambridge University Press, pp. 101-113.

Williamson, Timothy

(2000) *Knowledge and Its Limits*, Oxford: Oxford University Press.

(2014) "Very Improbable Knowing", *Synthese* 79: 971-999.

(2017) "Acting on Knowing", in J.A. Carter, E. Gordon, and B. Jarvis (eds.), *Knowledge-First*, Oxford University Press.

Worsnip, Alex

(2018) "The Conflict of Evidence and Coherence", *Philosophy and Phenomenological Research* 96.1: 3-44.