

Fragile Epistemic States and Rational Epistemic Akrasia

Sebastian Liu

Word Count: 8874

Abstract

Whether you're rational seems to depend both on whether your beliefs respect your evidence and whether your beliefs appropriately cohere with one another. But sometimes, when you're either rationally uncertain what your evidence is or what your evidence supports, conforming your beliefs to your evidence requires you to be epistemically akratic – to believe p while believing that p is unlikely on your evidence – a state which seems paradigmatically incoherent. In these instances, the putative rational requirements of believing in accordance with your evidence and avoiding incoherent combinations of beliefs appear to conflict. This paper aims to resolve this conflict by vindicating the rationality of epistemic akrasia. On the view that emerges, however, while it can be rational for you to be epistemically akratic, for principled reasons, you're always prohibited from believing that your evidence permits you to be akratic with respect to any particular proposition. Crucially, whenever your evidence licenses akratic beliefs, you'll be in an epistemic state which is 'fragile', and for principled reasons, you should never believe that you're in a fragile epistemic state. This characteristic of fragile epistemic states offers a natural and appealing explanation for why akratic beliefs appear irrational even when they're not, which has implications for peer disagreement, bias, knowledge iteration, and the relationship between substantive and structural rationality.

1 Evidence, Coherence, Epistemic Akrasia

Being rational is hard. The demands of rationality, it seems, include both evidential and coherence considerations:

(EVIDENTIALISM) Conform your beliefs to your evidence.¹

(COHERENCE) Avoid incoherent combinations of beliefs.²

According to EVIDENTIALISM you're required to believe p when your evidence adequately supports p , and to refrain from believing p when you lack any evidence for p . According to COHERENCE you're prohibited from simultaneously believing p and believing $\neg p$, and from believing that p is likely while believing that p is unlikely. Of course, even if your beliefs are coherent, it's often difficult to determine what your evidence supports when your evidence is complex or scarce, and so it's often difficult to abide by EVIDENTIALISM. And, even ignoring questions about what your evidence supports, it's often difficult to ensure that your beliefs are free of contradictions and probabilistically consistent when you have a sufficiently large number of beliefs, and so it's often difficult to abide by COHERENCE. Being rational is hard.

But it gets worse. For the putative requirements of EVIDENTIALISM and COHERENCE appear to conflict: In certain instances – when you're either rationally uncertain about what your evidence is or what your evidence supports – your evidence will license believing both p and that p is unlikely on your evidence. When you have such a combination of beliefs, you're epistemically akratic and you violate:

(ANTI-AKRASIA) Avoid the combination of believing p while believing that p is unlikely on your evidence.³

Akratic beliefs appear paradigmatically incoherent. For it seems clear that rationality does not permit you to have beliefs which you, from your own perspective, take to be unsupported

¹Those sympathetic to EVIDENTIALISM include Clifford (1897), Carnap (1950), Conee and Feldman (1985), Williamson (2000), Adler (2002), Kelly (2003), Feldman (2005), Shah (2006), White (2007), Greco (2014), Horowitz and Sliwa (2015), Salow (2019), and Lasonen-Aarnio (2020).

²For discussion of COHERENCE, see for instance Bonjour (1985), Kolodny (2005), Scanlon (2007), Broome (2013), Fitelson and Easwaran (2015), Kiesewetter (2017), Lord (2018), Pryor (2018), Worsnip (2018), and Fogal (2020).

³There are various formulations of ANTI-AKRASIA. See, for instance, Elga (2005, p.116), Feldman (2005, pp.108-110), Kolodny (2005, p.521), Christensen (2010a, pp.202-207), Huemer (2011, p.8), Smithies (2012, pp.280-285), Broome (2013, p.98), Greco (2014, pp.201-202), Horowitz (2014, p.718), White (2014, p.308), Horowitz and Sliwa (2015, pp.2847-2848), Titelbaum (2015, p.261), Littlejohn (2018, p.261), Worsnip (2018, p.13), Rinard (2019, p.249), Salow (2019, pp.399-400), and Lasonen-Aarnio (2020, p.600). Other candidate constraints related to ANTI-AKRASIA include the requirement that you avoid the combination of failing to believe p while believing that your evidence likely supports p , and the requirement that you avoid the combination of having credence C in p while thinking that the ideal credence in p given your evidence is C' , where $C \neq C'$. Though for expositional purposes my main focus in this paper will be on ANTI-AKRASIA, the coming discussion will also bear on these related putative requirements.

by your evidence. But if akratic beliefs are incoherent and there are cases in which your evidence supports believing akratically, then EVIDENTIALISM is in conflict with COHERENCE, as sometimes the requirement of respecting your evidence and the requirement of having structurally coherent beliefs diverge. Not only is it hard to be rational, but it seems that sometimes the demands of rationality themselves conflict.⁴

Only, they don't. That's my first claim in this paper. It's not that there aren't cases in which your evidence supports believing akratically. There *is* a conflict between EVIDENTIALISM and ANTI-AKRASIA. It's also not that EVIDENTIALISM and COHERENCE aren't genuine requirements of rationality. Rationality *does* require you to respect your evidence and have coherent beliefs. Rather, it's that COHERENCE doesn't entail ANTI-AKRASIA. Epistemic akrasia *isn't* always incoherent. Indeed it can be perfectly rational.⁵ But something in the vicinity of ANTI-AKRASIA is true. My second claim is that what rationality demands is:

(ANTI-AKRASIA*) Avoid believing that your evidence permits you to believe both p and that p is unlikely on your evidence.

ANTI-AKRASIA* does not prohibit you from being epistemically akratic. It does, however, prohibit you from believing that your evidence permits you to be epistemically akratic with respect to any particular proposition. Importantly ANTI-AKRASIA* is consistent with both EVIDENTIALISM and COHERENCE. Properly understood, ANTI-AKRASIA* explains why violations of ANTI-AKRASIA appear (and merely appear) irrational. The central aim of this paper is to defend ANTI-AKRASIA* and argue that it accommodates the considerations that purportedly favor ANTI-AKRASIA.

2 A Conflict

To start, it will be helpful to make concrete the conflict between EVIDENTIALISM and ANTI-AKRASIA: Believing in accordance with your evidence can sometimes require you to believe both p and that p is unlikely on your evidence. The next two subsections illustrate this with two examples, one involving an instance of access failure, in which you're uncertain what your evidence is, and one involving an instance of self-misleading evidence, in which you're uncertain what your evidence supports. The third subsection discusses the conflict more generally and makes explicit some assumptions of this paper.

⁴Couldn't there be instances of 'rational incoherence', perhaps in preface-style cases? If so, COHERENCE would only require avoiding certain *kinds* of incoherent combinations of beliefs. I want to largely set aside the question of whether rational incoherence is possible; if you think that it is possible, this paper can be read as defending the claim that epistemic akrasia, even if it is incoherent, can be perfectly rational. For discussion, see for example Kolodny (2007), Broome (2013), and Kiesewetter (2017).

⁵Others rejecting or expressing doubt (to varying degrees) about ANTI-AKRASIA include for instance Coates (2012), Hazlett (2012), Wedgwood (2012), Williamson (2014), Dorst (2019), Weatherson (2019), Lasonen-Aarnio (2020), and Hawthorne, Isaacs, and Lasonen-Aarnio (2021). My approach here, however, is importantly different and doesn't, for instance, rely on the strategy of distinguishing between different evaluative perspectives such as competence and success.

2.1 Failures of Access

Does your evidence always entail what your evidence is? That is, if E is your total evidence, does E always entail that E is your total evidence? It seems not, for it seems that sometimes, you can be rationally required to be uncertain about what your evidence is. Such cases involve a failure of access.⁶ Consider:

(STUDENT) You're a student taking an important history exam. One of the questions asks you for the year that William the Conqueror landed in England. As it happens, you have learned many facts about William the Conqueror from class and from reading your textbook, including this specific year. But while the correct answer, 1066, comes to your mind, because of your dispositions, you think that your answer is very likely just a guess out of desperation. Due to the pressure and stress you're experiencing, you don't remember learning that William the Conqueror landed in England in 1066. You recognize that as far as you can tell, you have no evidence for this answer, for as far as you can tell, you are just guessing.⁷

Here's an attractive (although contestable) diagnosis of STUDENT. You have good, even conclusive, evidence for your answer, 1066. After all, you have properly learned it, can reliably recall it, and (at least on some accounts of knowledge) know it. But despite having evidence for your answer, you nevertheless lack evidence that you have evidence for your answer. You have no recollection of learning the year so for all you know, you're merely guessing.

If your evidence supports that William the Conqueror landed in England in 1066, but you also lack evidence that you have this evidence, then STUDENT is an example in which conforming your beliefs to your evidence requires you to violate ANTI-AKRASIA, since conforming your beliefs to your evidence requires you to believe both that William the Conqueror landed in England in 1066, and also that this is likely not supported by your evidence. STUDENT is an example of uncertainty about what your evidence is as a result of access failure, and respecting your evidence in light of such uncertainty leads to epistemic akrasia.

2.2 Self-Misleading Evidence

Does your evidence always entail what your evidence supports? That is, if your total evidence E supports p , does E always entail that E supports p ? It seems not, for it seems that

⁶More precisely, these would be (alleged) failures of 'positive access', the thesis that whenever you have a piece of evidence, your evidence entails that you have that piece of evidence. There are also (alleged) failures of 'negative access', the thesis that whenever you lack a piece of evidence, your evidence entails that you lack that piece of evidence. See footnote 12 for an example.

⁷This case comes from Radford (1966). Examples like this are often taken to be counterexamples to KK, the thesis that knowledge entails knowledge of knowledge. If evidence is equated with knowledge, failures of (positive) access just are failures of KK. See also, for instance, Feldman (2005) and Worsnip (2018) for discussion about how these kinds of cases relate to epistemic akrasia.

sometimes, you can be rationally required to be uncertain about what your evidence supports. Such cases are ones in which what your evidence supports and what your evidence supports about what it supports diverge.⁸ Consider:

(HYPOXIA) You're piloting a small aircraft and you're currently a few hundred miles from your intended destination. The fuel gauge looks dangerously close to empty. You wonder whether you will need to make an emergency landing, so you perform a series of intricate calculations based on the information provided by the flight instruments. You conclude that you have enough fuel after all.

You then receive a message from ground control informing you that the plane's current altitude makes you highly susceptible to hypoxia. If you are suffering from hypoxia, then although you will not be able to tell, your reasoning abilities will be severely impaired.

Unknown to you, you are lucky – you have not been affected by hypoxia and you have in fact perfectly evaluated the evidence from the flight system.⁹

Here's an attractive (although contestable) diagnosis of HYPOXIA. You have strong evidence that you have enough fuel to safely arrive at your destination. The information you have about your current mileage, altitude, and speed, and your impeccable calculations based on this information suggest that you don't need to make an emergency landing. But you also have strong, albeit misleading evidence for thinking that your evidence likely doesn't support your conclusion that you have enough fuel. If it's probable that you're hypoxic, then you should suspect that you've failed to properly appreciate your flight system evidence which bears on the question of whether you have enough fuel, for despite being misled by ground control, you have reason to think that you've made a mistake in your calculations.

Yet if your evidence supports both that you have enough fuel and also that this is unlikely on your evidence, then in HYPOXIA respecting your evidence requires you to believe akratically. HYPOXIA is an example in which you're uncertain about what the evidence supports because the evidence is self-misleading, and such uncertainty mandates violating ANTI-AKRASIA.

2.3 The Conflict More Generally

When you lack access to your evidence as in the unconfident student example, you're uncertain what your evidence is. When your evidence is self-misleading as in the hypoxic pilot example, you're uncertain what your evidence supports. STUDENT and HYPOXIA, however,

⁸Your total evidence E may also be self-misleading if there are cases in which E fails to support p but E fails to support that E fails to support p .

⁹This case is widely discussed. See for example Christensen (2010b), Elga (2013), Schechter (2013), Lasonen-Aarnio (2014), Schoenfield (2015), and Neta (2019).

are simply instances of a much more general conflict between EVIDENTIALISM and ANTI-AKRASIA, a conflict which doesn't depend on the specific details of these two cases.

First, the solicited judgments about STUDENT and HYPOXIA can be resisted. For example, you may be inclined to think that the fact that you have a belief about when William the Conqueror landed in England is itself evidence for your answer in STUDENT. Or, in HYPOXIA, you might think that perhaps upon receiving the warning from ground control, information from your flight system should cease to be considered as part of your evidence. The status of these more controversial points, however, is largely insignificant, for unless you can *never* be rationally uncertain about what your evidence is or what your evidence supports, there will be instances in which conforming your beliefs to your evidence requires some degree of epistemic akrasia.¹⁰ For this reason, and because it's helpful to have concrete examples for expositional purposes, this paper will take for granted that STUDENT and HYPOXIA are cases in which complying with EVIDENTIALISM requires violating ANTI-AKRASIA.¹¹

Second, the tension between EVIDENTIALISM and ANTI-AKRASIA isn't meant to depend on any particular assumptions about the nature of evidence. However, in this paper, evidence will be understood in a broadly 'externalist' sense: Your evidence will be taken to consist of a set of true propositions, which can include propositions about the external world. This assumption about evidence doesn't immediately resolve the issue of interest. On the contrary, it's often thought to be an especially unwelcome consequence of externalist accounts of evidence that the access principles become indefensible, and failures of access lead to violations of ANTI-AKRASIA (given EVIDENTIALISM).¹² If you're already committed to externalism about evidence, this paper can be read as a defense of EVIDENTIALISM in light of its tension with ANTI-AKRASIA. If you're more sympathetic to 'internalist' accounts of evidence,

¹⁰For discussion see especially Dorst (2019) and Worsnip (2019); for possible lines of resistance, see for instance Greco (2019), Neta (2019), Skipper (2019), and Smithies (2019).

¹¹I've presented STUDENT and HYPOXIA in terms of full, or outright belief. But the analysis applies more generally to credences as well: In both examples, your total evidence arguably supports having a high credence C in some proposition p , while also supports thinking that the ideal credence to have in p given your evidence is lower than C .

¹²Some may be tempted to defend qualified forms of the access principles. For simplicity, I'll be ignoring those details, though for possible difficulties, see in particular San (2019) and Liu (2020). While I'll focus on cases of positive access failure, the kind of strategy I pursue is straightforwardly applicable to cases of negative access failure as well. Let me take this opportunity to illustrate a violation of negative access and discuss how it bears on a related (putative) ANTI-AKRASIA requirement mentioned in footnote 3, according to which rationality requires you to avoid having credence C in p while thinking that the ideal credence in p on your evidence is C' . Suppose you're facing a clock which reads reads 12:00. In fact, the clock is broken, and it's currently 1:00. Given that your evidence consists of a set of true propositions, your evidence won't include p , the proposition that it *is* 12:00, though it will include q , the proposition that the clock *reads* 12:00. Yet, assuming you have no reason to think that the clock is broken, your evidence presumably won't include the proposition that your evidence doesn't include p . Suppose that, given q , you should have 0.9 credence in p . Then, while you should have 0.9 credence that it's 12:00, you should also estimate that the ideal credence in p is $0.9 \times 1 + 0.1 \times 0.9 = 0.99$. So your credence that it's 12:00 should be 0.9, but you should think that your 0.9 credence is too low, for you should think that the ideal credence to have in p is 0.99. For further discussion see especially Silins (2005), White (2014), and Salow (2019).

this paper can be read as arguing in favor of externalizing evidence to the extent that it can resolve the conflict between EVIDENTIALISM and ANTI-AKRASIA.¹³

3 The Alleged Irrationality of Epistemic Akrasia

So: The requirements of respecting your evidence and avoiding epistemic akrasia can come into conflict in cases involving failures of access or self-misleading evidence. There are a number of potential resolutions, the most obvious ones being to reject EVIDENTIALISM, reject ANTI-AKRASIA, or conclude that it is sometimes impossible to satisfy all the demands of rationality. None of these options appears attractive, but maintaining that epistemic akrasia can be rational seems particularly egregious. ANTI-AKRASIA enjoys significant intuitive appeal. It's a widely shared sentiment that rationality requires your beliefs to be aligned with your beliefs about what you should believe, or what your evidence supports:

Irrationality in the clearest sense occurs when a person's attitude fails to conform to his or her own judgments: when, for example, a person continues to believe something... even though he or she judges there to be good reason for rejecting it... (Scanlon 1998, p.25)

If you doubt that [ANTI-AKRASIA] is a norm, you can make its plausibility vivid by imagining an argument with an opponent who violates it. You marshal your best evidence for your view. Your opponent agrees that you've presented strong evidence for your view, and has no counter-evidence. But no matter how much evidence you present, or how strong it is, he gains no confidence in your view... in this infuriating scenario, your opponent is being unreasonable. (Elga 2005, p.116)

It is irrational to believe either (i) the proposition that p , or (ii) the proposition that one has justification to believe that p , while disbelieving or withholding belief in the other. In other words, believing a proposition rationally commits one to believing that one has justification to believe it and vice versa. (Smithies 2012, p.284)

The apparent unpalatability of rejecting ANTI-AKRASIA is not just limited to intuitive judgments. To allow for the possibility of rational epistemic akrasia also requires answering a number of challenges arising from assertion, action, and reasoning.¹⁴

¹³On internalist accounts of evidence, your evidence consists of, for example, your experiences, or your sense data. But even if internalism about evidence can preserve the access principles – even if on internalist accounts of evidence, you're never rationally uncertain what your evidence is – your evidence may still require you to believe akratically, for you may still have self-misleading evidence.

¹⁴These arguments are discussed by, among various others, Feldman (2005), Christensen (2010a), Smithies (2012), Elga (2013), Greco (2014), Horowitz (2014), Horowitz and Sliwa (2015), and Worsnip (2018).

Assertion. If you can be rationally akratic, then it seems that you should be able to express your beliefs by making assertions of the form ‘ p , although p is unlikely on my evidence’, or ‘while my evidence likely does not support p , I believe p ’, or ‘I am confident that p , and I am also confident that p is improbable given my evidence’. These kinds of assertions appear infelicitous and plausibly shouldn’t be asserted, but it’s unclear why they would be impermissible to assert if epistemic akrasia could be rational.

Action. Epistemic akrasia licenses a form of practical akrasia, which involves performing some action while judging that you likely have most reason not to perform that action. Practical akrasia is commonly taken to be a paradigm of practical irrationality. But if you’re epistemically akratic, then it seems you should be willing to bet on p , since you think that p is true, and also willing to bet on the proposition that you’ll likely lose your bet on p , since you think that p is unlikely on your evidence. Accepting this pair of bets is manifestly bizarre, because you would be accepting at least one bet which you do not, in some sense, fully endorse. If practical akrasia is irrational, there’s significant pressure to think that epistemic akrasia is irrational as well.

Reasoning. If your akratic beliefs were rational, you can presumably reason from ‘my evidence likely does not support p , but p is true’ to either ‘I lack access to the fact that I have evidence for p ’ or ‘my evidence is self-misleading with respect to p ’. Yet this pattern of reasoning is surely inappropriate, for it seems to be in the nature of cases that involve failures of access and self-misleading evidence that you’re not in a position to determine such facts about access and misleading evidence from only the structural features of your beliefs.

These formidable challenges to the rationality of epistemic akrasia are strong reasons for endorsing ANTI-AKRASIA.¹⁵ If it’s irrational to be epistemically akratic, then it’s unsurprising why asserting, acting on, or reasoning from akratic beliefs would appear problematic: You would be asserting, acting on, or reasoning from combinations of beliefs that you should not have.

The following two sections argue that these objections, despite what they initially appear to suggest, don’t impugn the rationality of akratic beliefs themselves. Crucially, once the question of what you should believe and the question of what you should believe about what your evidence supports are separated, these objections can be accommodated by ANTI-AKRASIA*, for ANTI-AKRASIA* is sufficient for explaining why akratic beliefs, even if they are rational, will appear irrational when you assert, act on, or reason from them. The next section argues for this in cases involving failures of access. The section after relates cases of access failure to cases of self-misleading evidence.

¹⁵Unsurprisingly, in light of these considerations, it’s fairly uncommon to reject ANTI-AKRASIA. For a sampling of other strategies, see Elga (2013), Horowitz (2014), Titelbaum (2015), Littlejohn (2018), Neta (2018), Worsnip (2018), and Salow (2019).

4 Failures of Access and Unknowability

4.1 The Structure of Failures of Access

What is the structure of your epistemic state in STUDENT? Let F be the general factive relation – for instance, knowledge, or what you’re in a position to know – such that whenever you stand in F to a proposition, that proposition is part of your evidence. Failures of access in cases like STUDENT then most naturally correspond to the epistemic state $Fp \wedge \neg FFp$, since if p (that William the Conqueror landed in England in 1066) is part of your evidence, then Fp , and if you lack evidence that your evidence includes p , then $\neg FFp$. Call an epistemic state ‘fragile’ just in case it is impossible for you to have conclusive evidence or know that you are in such a state.¹⁶ Observation: Failures of access are one kind of fragile epistemic state. For suppose your evidence conclusively supports that you’re in a particular case of access failure. Then your evidence includes both Fp and $\neg FFp$. Since p is part of your evidence if and only if Fp , it follows that FFp and $F\neg FFp$. But F is factive, so $F\neg FFp$ entails $\neg FFp$, and therefore FFp and $\neg FFp$. Contradiction; you can’t have conclusive evidence that you’re in a particular case of access failure.¹⁷ An analogous argument establishes that it’s impossible for you to *know* that you are in a particular case of access failure.¹⁸ Because (un)knowability is a more familiar concept than the (im)possibility of having conclusive evidence, the subsequent discussion focuses on the consequences of knowability as opposed to the possibility of having conclusive evidence. The arguments, however, are fully general.

4.2 A Constraint on Belief

Since whenever you’re in a particular case of access failure, you’re in a fragile epistemic state, when you’re in a particular case of access failure, you won’t be in a position to know that you are, and so in these instances, you won’t be in a position to know that your evidence permits

¹⁶The proposition that you are in a particular fragile epistemic state is an example of a ‘blindspot’ proposition in the sense of Sorensen (1988).

¹⁷This observation has been made by, among others, Williamson (2000) and Worsnip (2018). Indeed, the formal assumptions required are weaker than those made in the main text. Assume in the background a bimodal logic with two modal operators, F and \Box , where Fp iff p is part of your evidence and $\Box p$ iff p is metaphysically necessary, and abbreviate $\neg\Box\neg$ as \Diamond . Then it’s easy to show $\neg\Diamond F(Fp \wedge \neg FFp)$ given:

$$\text{(FACTIVITY)} \quad Fp \rightarrow p$$

$$\text{(DISTRIBUTION)} \quad F(p \wedge q) \rightarrow \Diamond(Fp \wedge Fq)$$

$$\text{(NECESSITATION)} \quad p/\Box p$$

It’s also easy to show $\neg\Diamond F(\neg Fp \wedge \neg F\neg Fp)$, which corresponds to failures of negative access. Consequently, assuming that F is factive and satisfies a weak distribution principle, and that \Box obeys necessitation, you can never have conclusive evidence (or know) that you’re in a particular case of positive or negative access failure, so the kinds of cases discussed in footnote 12 can be given the same treatment.

¹⁸It’s easy to see that, where K is the knowledge operator and your epistemic state in STUDENT is understood to be $Fp \wedge \neg KFp$, then $\neg\Diamond K(Fp \wedge \neg KFp)$ given that K is factive and distributes over conjunction.

believing akratically when it does. This subsection defends the claim that, because being in a fragile epistemic state is unknowable, you should never believe that you're presently in such a state.¹⁹ The next subsection argues that this claim supports ANTI-AKRASIA*.

It is impermissible for you to believe that you're in a particular fragile epistemic state if the following putative norm constrains belief:

(KNOWABILITY-BELIEF-NORM) Believe p only if p is knowable.²⁰

KNOWABILITY-BELIEF-NORM is deeply plausible. On an impressionistic level, if belief functions to eliminate those possibilities incompatible with the believed proposition, then arguably you shouldn't believe unknowable propositions because you're not in a position to eliminate those possibilities. More concretely, deniers of KNOWABILITY-BELIEF-NORM face at least two challenges.

First challenge. Consider Moorean conjunctions such as:

1. It's raining and I don't believe that it's raining.
2. I don't know that there are two hands in front of me, but there are two hands in front of me.
3. I'm not dreaming right now, although it's an open question whether I am.
4. I lack justification to believe that the animal in the zoo pen is a zebra, but it's a zebra.

It seems clear that it is infelicitous to assert (1) – (4). The observation that Moorean assertions are infelicitous is often taken to support the knowledge norm of assertion, according to which you ought assert only what you know.²¹ Since Moorean conjunctions are unknowable, you ought not assert them. Given the further reasonable suggestion that belief is the inner analogue of assertion, it is natural to think that there is a corresponding knowledge norm of belief – you ought believe only what you know.²² If knowledge is the norm of belief, then KNOWABILITY-BELIEF-NORM immediately follows since you can't know what's unknowable.

Of course, you may doubt that knowledge is the norm of assertion and belief. Even so, you should still accept KNOWABILITY-BELIEF-NORM. Because Moorean conjunctions like (1) – (4) are unknowable, to prohibit Moorean assertions, it is sufficient to prohibit asserting

¹⁹In fact, being in a case of access failure is *knowably unknowable* and the arguments in this subsection can be modified to rely only on this weaker notion. It's somewhat cumbersome to talk about knowable unknowability, so I'll focus just on unknowability.

²⁰On some accounts of belief, such as the one advanced by Stalnaker (2006), you believe p just in case you don't know that you don't know p . Since, if p is unknowable, you don't know p , on these accounts of belief, KNOWABILITY-BELIEF-NORM is trivial.

²¹For endorsement of the knowledge norm of assertion, see, among others, Unger (1975), Williamson (2000), Adler (2002), Sutton (2005), Bird (2007), and Huemer (2007).

²²See for instance Unger (1975), Williamson (2000), DeRose (2002), and Hawthorne (2004). For my purposes, the weaker norm that you ought believe only what you're in a position to know is already sufficient.

unknowable propositions. But if you ought not assert that which is unknowable, then conditional on the claim that what shouldn't be asserted also shouldn't be believed, you ought not believe that which is unknowable. And this is just what KNOWABILITY-BELIEF-NORM requires.

The first challenge for a denier of KNOWABILITY-BELIEF-NORM is to explain why Moorean conjunctions should not be believed.

Second challenge. Consider lottery propositions. In a fair (finite) lottery, it seems that while you should be highly confident (proportional to the number of tickets in the lottery) that any random ticket will lose, you shouldn't outright *believe* that it will lose. After all, if you believe that ticket 1 will lose, and you also believe that ticket 2 will lose, won't you – perhaps after a moment of reflection – believe that both tickets 1 and 2 will lose? This quickly leads to absurdity, for if, for any given any ticket t you believe that t will lose, then (assuming your beliefs are closed under conjunction) you will also believe that all the tickets will lose. But you know (and therefore believe) that some ticket in the lottery will win, so you will have inconsistent beliefs. One inviting explanation for why you shouldn't believe lottery propositions is that belief requires credence 1.²³ Although your credence that t will lose should be quite high, it surely shouldn't be 1, and so you shouldn't believe that t will lose. If belief requires credence 1, then presumably KNOWABILITY-BELIEF-NORM follows because credence 1 corresponds to subjective certainty and you shouldn't be subjectively certain in unknowable propositions.

But you may be reluctant to accept that belief requires maximal credence. Nevertheless, you should still accept KNOWABILITY-BELIEF-NORM. On standard views, you can't know lottery propositions.²⁴ Provided that lottery propositions are unknowable, KNOWABILITY-BELIEF-NORM offers an independent explanation for why you shouldn't believe that some particular ticket will lose: You shouldn't believe lottery propositions because you shouldn't believe unknowable propositions.

The second challenge for the denier of KNOWABILITY-BELIEF-NORM is to explain why lottery propositions should not be believed.

So: You shouldn't believe p if p is unknowable. KNOWABILITY-BELIEF-NORM is supported by considerations from Moorean conjunctions and lottery propositions. Because it's

²³Those who endorse that belief requires credence 1 include Clarke (2013), Greco (2015a), and Dodd (2017). There are two potential complications that I won't be able to address here. The first concerns arguments which purport to show, by appeal to linguistic data, that belief is weak, perhaps merely requiring thinking that p is likely. See in particular Hawthorne, Rothschild, and Spectre (2016). The second concerns the straightforward denial that you shouldn't believe lottery propositions. See especially Kyburg (1961) on accepting that you should believe lottery propositions and rejecting that rational belief must be closed under conjunction.

²⁴See for example Lewis (1996), Nelkin (2000), and Hawthorne (2004).

impermissible to believe what's unknowable, it's impermissible for you to believe that you're in a particular fragile epistemic state.²⁵

4.3 The Inevitable Appearance of Irrationality

ANTI-AKRASIA* forbids you from believing that your evidence permits you to believe both p and that p is unlikely on your evidence. Because it's impermissible for you to believe that you're in a particular fragile epistemic state and failures of access are fragile epistemic states, it's impermissible for you to believe that you're in a particular case of access failure. It follows that, in cases of access failure, it's impermissible for you to believe that your evidence supports believing akratically.²⁶ And these are exactly the kind of beliefs that ANTI-AKRASIA* forbids. Rationality requires compliance with ANTI-AKRASIA*.

That ANTI-AKRASIA* is a requirement of rationality does not imply that ANTI-AKRASIA is not. ANTI-AKRASIA* prohibits a certain kind of belief about what your evidence supports, while ANTI-AKRASIA prohibits a certain combination of beliefs. Why should accepting ANTI-AKRASIA* bear on whether epistemic akrasia can be rational? The answer is that the considerations which motivate ANTI-AKRASIA can be captured by ANTI-AKRASIA*. That is, while ANTI-AKRASIA explains the intuitive irrationality of epistemic akrasia and the apparent problems arising from asserting, acting on, and reasoning from such beliefs, so does ANTI-AKRASIA*. The difference is that ANTI-AKRASIA* has the resources to explain why epistemic akrasia appears irrational without attributing irrationality to the combinations of beliefs themselves. Here is why.

Fragile epistemic states are unstable under reflection. Asserting, acting on, and reasoning from your akratic beliefs will tend to make salient the fact that you have those akratic beliefs. When you come to recognize that you're epistemically akratic, it becomes difficult to comply with ANTI-AKRASIA*. What should you think about your beliefs when you come to recognize that you're akratic? If you maintain that your akratic beliefs are rational because they are supported by your evidence in a case of access failure, then you would be violating ANTI-AKRASIA*. In these circumstances, it is only natural for you to think that not all of your beliefs are properly proportioned to your evidence, even when they are.

This point deserves emphasis. Typically when your beliefs respect your evidence, it's unproblematic for you to have the additional belief – which tends to accompany your beliefs when they are made salient – that you're respecting your evidence.²⁷ But when you have

²⁵I want to set aside cases generated by Frege puzzles, in which you're unsure about who you are. For example, suppose Clark Kent believes that he isn't Superman. Then perhaps it's permissible for him to believe both that it's raining and that Superman doesn't believe that it's raining, even though this proposition is unknowable for Superman (and therefore unknowable for Clark Kent). These kinds of examples are orthogonal to the paradigm cases of access failure and self-misleading evidence.

²⁶My proposal is similar in spirit to the one in Salow (2019), but differs importantly because my proposal does not rely on contextualism about evidence.

²⁷See for example Adler (2002, p.26) and Worsnip (2021, p.132) for related remarks.

akratic beliefs in cases of access failure, regardless of whether those beliefs are supported by your evidence, it's impermissible for you to have this additional belief given ANTI-AKRASIA*. In these cases, maintaining that your evidence licenses believing akratically is tantamount to violating this requirement. Yet if, in compliance with ANTI-AKRASIA*, you refrain from believing that your evidence supports your akratic beliefs even when these beliefs are made salient, there will be rational pressure for you to think that your beliefs fail to respect your evidence and that you're irrational. This is characteristic of fragile epistemic states: Once you come to recognize that you're in such a state it'll tend to seem to you that you're irrational.

That's the story so far. When you are epistemically akratic as a consequence of conforming your beliefs to your evidence in cases of access failure, it'll appear to you that you're irrational when you come to recognize that you hold those beliefs. ANTI-AKRASIA*, which follows from KNOWABILITY-BELIEF-NORM and the structure of fragile epistemic states, predicts this phenomenon, so at least in cases of access failure, there's no need to appeal to ANTI-AKRASIA – there's no need to insist that it's the akratic beliefs themselves that are irrational – to explain why asserting, acting on, or reasoning from such beliefs appear problematic.

5 Self-Misleading Evidence from Failures of Access

But that isn't the full story. Even if ANTI-AKRASIA* adequately explains why epistemic akrasia appears (and merely appears) irrational in cases of access failure like STUDENT, it isn't obvious whether, or to what extent, such an explanation is relevant to cases of self-misleading evidence like HYPOXIA. The appeal to ANTI-AKRASIA* crucially relies on the fact that when you're uncertain what your evidence is, you're in a fragile epistemic state, but it doesn't seem that when you're uncertain what your evidence supports, you're also in a fragile epistemic state. What is the relationship between cases of access failure and cases of self-misleading evidence?

A tempting thought is that these are simply different and distinct ways in which proportioning your beliefs to your evidence requires you to be epistemically akratic:²⁸

... it's important to distinguish two quite different ways to be misled about what one's evidence supports. On one hand, there is the possibility that one might be misled... about what one's evidence *is*. On the other hand, there is the possibility that one is misled... about the *evidential support relations*... (Worsnip 2018, p.19)

There are two sorts of reason I might be uncertain about what my evidence supports. First, while I might know what's supported by each possible body of evidence, I might fail to know which body of evidence is *mine*. Second, while

²⁸In addition to those quoted below, others, including Elga (2013, p.132), Titelbaum (2015, p.262), Worsnip (2018, p.19), Lasonen-Aarnio (2019, pp.156-157), Salow (2019, p.402), and Hedden and Dorst (2022, p.421) also seem to sharply separate cases of access failure and cases of self-misleading evidence.

I might know which body of evidence is mine, I might fail to know what that given body of evidence supports. (Greco 2019, pp.85-86)

The thought that failures of access and self-misleading evidence are distinct ways in which proportioning your beliefs to your evidence requires you to be epistemically akratic is mistaken. For suppose you're unsure whether p is part of your evidence. Suppose further that you know that if p is part of your evidence, your evidence would strongly support q (imagine that you know p entails q), and that if p is not part of your evidence, your evidence would strongly support $\neg q$. Here, uncertainty about what your evidence is directly bears on uncertainty about what your evidence supports, for it is a consequence of your uncertainty about whether p is part of your evidence that you are uncertain whether your evidence supports q .²⁹

Since uncertainty about what your evidence is can influence (indeed can be the source of) uncertainty about what your evidence supports, examples like HYPOXIA should not be assumed to be instances of uncertainty about evidential support as opposed to instances of uncertainty about evidence. Why think that when you have reason to suspect that you're susceptible to hypoxia, you should become unsure what the probative force of your evidence is, as opposed to becoming unsure which propositions are part of your evidence? Consider:

(HYPOXIA*) You're piloting a small aircraft and you're currently a few hundred miles from your intended destination. The fuel gauge looks dangerously close to empty. You wonder whether you will need to make an emergency landing, so you perform a series of intricate calculations based on the information provided by the flight instruments. You conclude that you have enough fuel after all.

You then recall from your training that at certain altitudes pilots become susceptible to hypoxia. If you are hypoxic, then although you will not be able to tell, your reasoning abilities will be severely impaired. Because you cannot recall exactly which altitudes are dangerous, you contact ground control to ask. You are informed that your current altitude is safe.

In fact you're not hypoxic – you have enough fuel, and you've evaluated the evidence from the flight system perfectly.

Consider, for instance, the proposition that you have enough fuel. Is this proposition part of your evidence in HYPOXIA*? Your fuel calculations are impeccable, and your perceptual and computational faculties are functioning reliably. If your evidence includes propositions like that you are 300 miles from your destination, that the aircraft is flying at 150 miles per hour, or that your current altitude does not make you susceptible to hypoxia, then it would be odd to disqualify the proposition that you have enough fuel from being included as part

²⁹I don't mean to suggest that uncertainty about what your evidence is always implies uncertainty about what your evidence supports. If you know that q is part of your evidence, then you will trivially know that your evidence supports q , even if you are uncertain whether p is part of your evidence.

of your evidence. Suppose ground control asks you why you did not make an emergency landing. In normal circumstances, it is perfectly natural and acceptable for you to cite the fact that you had enough fuel to safely continue flying in response. In HYPOXIA*, it seems plausible that your evidence includes the proposition that you have enough fuel.

Notice that HYPOXIA and HYPOXIA* differ not with respect to whether you're hypoxic, but with respect to whether ground control informs you that you're susceptible to hypoxia. In HYPOXIA, you have reason to think that your faculties have been severely impaired, and in HYPOXIA*, you don't. Prior to receiving the warning from ground control, it's reasonable to think that the proposition that you have enough fuel is part of your evidence in *both* cases, for prior to receiving the warning, HYPOXIA and HYPOXIA* are the same. Does the warning from ground control preclude this proposition from being (or remaining) part of your evidence in HYPOXIA? It's not obvious that it does. By stipulation, you're not any less competent or reliable at evaluating your evidence, for you're not in fact hypoxic. A promising suggestion is that the hypoxia warning merely *limits your access* to your evidence: Once you have sufficient reason to suspect that you are hypoxic, propositions like that you have enough fuel (or that it's *likely* that you have enough fuel, or that you'll safely arrive at your destination) are no longer accessible to you, though still remain as part of your evidence.

Examples like HYPOXIA are ones in which you're unsure what your evidence *supports*. The present proposal is that they are *also* ones in which you're unsure what your evidence *is*. Indeed, it's in virtue of the uncertainty about what your evidence is that you're uncertain about what your evidence supports. While your evidence in HYPOXIA includes propositions like that you have enough fuel, the message from ground control is evidence that you lack reasons for thinking that these propositions are part of your evidence.³⁰ And plausibly, when you have sufficiently strong reason to doubt that you have some piece of evidence which you in fact have, you'll be in a case of access failure. But failures of access are not failures of rationality; that your evidence does not always entail what your evidence is, is no fault of your own and does not reflect some rational defect on your part.

Understanding examples of uncertainty about evidential support as involving failures of access unifies cases like STUDENT and HYPOXIA, and moreover, it also preserves the guiding intuitions about the alleged irrationality of violating ANTI-AKRASIA. If cases of self-misleading evidence involve failures of access, then insofar as ANTI-AKRASIA* satisfactorily explains why epistemic akrasia (merely) appears irrational in cases like STUDENT, it also satisfactorily explains why epistemic akrasia (merely) appears irrational in cases like HYPOXIA. Just like in STUDENT, in HYPOXIA, if you proportion your beliefs to your evidence, you will

³⁰In fact, the warning from ground control presumably isn't just evidence that you don't currently stand in relation *F* to the proposition that you have enough fuel, but also evidence that you're not in a position to stand in relation *F* to that proposition. So given that, once you have the evidence afforded to you by your flight instruments, you're in a position to stand in relation *F* to the proposition that you have enough fuel, your epistemic state would be $Pp \wedge \neg Fp$, where you stand in relation *P* to *p* just in case you're in a position to stand in *F* to *p*. It's straightforward to show that $Fp \wedge \neg Fp$ is also a fragile epistemic state – that $\neg \Diamond(Fp \wedge \neg Fp)$ – given natural assumptions.

be epistemically akratic. And crucially, just like in *STUDENT*, in *HYPOXIA*, having akratic beliefs doesn't entail that you're irrational. If *HYPOXIA* involves a failure of access, you'll be in a fragile epistemic state, so it will seem to you that your akratic beliefs are irrational when you assert, act on, or reason from them. *ANTI-AKRASIA** can account for why epistemic akrasia inevitably and predictably *seems* problematic while still allowing that it *can* be rational for you to believe both that you have enough fuel and also that this is unlikely on your evidence. What you're prohibited from believing – for principled reasons – is that your evidence permits having this combination of beliefs.

6 Disagreement, Bias, KK, Structure

The aim of previous two sections has been to argue that *ANTI-AKRASIA** can satisfactorily explain why akratic beliefs appear irrational even if they aren't, and so there is no additional reason to appeal to *ANTI-AKRASIA* to account for the problems arising from asserting, acting on, and reasoning from akratic beliefs. This section highlights several noteworthy consequences of the general picture here.

Peer Disagreement. Suppose you and I are roughly equally capable at evaluating the evidence that bears on some proposition p . Upon assessing the evidence you come to believe p and I come to believe $\neg p$. When we learn of our disagreement, how, if at all, should we revise our beliefs?³¹ The fact that we disagree is some evidence; in particular, it's evidence that at least one of us has failed to properly appreciate our evidence. It seems then, that when I learn of our disagreement, my evidence – which includes the fact that we disagree – will support believing akratically. Consider:

(DETECTIVE) You and I are equally competent detectives tasked with solving a case involving the theft of an expensive necklace. After examining the sets of fingerprints, the fresh mud tracks near the mansion, and the testimony of various suspects, I come to believe that the butler is guilty while you come to believe that the maid is guilty. When we convene, we learn of our disagreement.

As it so happens, the evidence actually supports that the butler stole the necklace. But the fact that you, someone who is as capable a detective as I am, came to a different conclusion after evaluating the same evidence is reason for me to think that I've made a mistake and that the evidence likely does not support my conclusion.

What should I believe in light of our disagreement? If I maintain my belief that the butler is guilty, then I would be ignoring the pertinent fact of our disagreement. And if I suspend judgment about who is guilty, then I would be ignoring the evidence which points to the

³¹On peer disagreement, see especially Kelly (2005), Christensen (2007), Elga (2007), and Hawthorne and Srinivasan (2013).

actual culprit. But to believe both that the butler is guilty and also that this is unlikely on my evidence just seems absurd. Not so, if epistemic akrasia can be rational in peer disagreement cases. Disagreement can be understood as functioning to prevent my access to certain propositions, such as that the butler is guilty or that the butler likely committed the crime, propositions which are part of my evidence when I've properly assessed the relevant clues. When imagining the case, it should not be assumed that I will have the further belief that my combination of beliefs is permitted by my evidence if I believe akratically. If epistemic akrasia can be rational, then even upon disagreeing in *DETECTIVE*, it can be rational for me to believe that the butler is guilty while also believing that this is unlikely on my evidence. Moreover, the judgment that there's something strange about my blatantly holding steadfast to my beliefs despite admitting that the disagreement should have some evidential force can still be captured: To remain steadfast while acknowledging our peer disagreement will tend to require me to believe that my combination of beliefs is supported by my evidence – which is impermissible given *ANTI-AKRASIA**.

Distorting Bias. It is a well-established empirical fact that we have a tendency to overrate ourselves – we often think we're better, smarter, more charismatic than we actually are.³² How should we take this fact into account when we're evaluating beliefs about ourselves? Consider:

(*POPULARITY*) I am wondering whether I am popular among my friends and colleagues. I gather evidence in an attempt to uncover what others think of me. I send out anonymous questionnaires asking for opinions from my friends and colleagues and interview a random sample of them. Upon assessing the responses, I become highly confident that I am very popular.

I am happy about this discovery. But I soon realize that because I have a strong preference to be popular, it's plausible that I have been biased in my assessment. In particular, it's likely that I have been giving more weight to positive feedback and less weight to negative feedback. Because it is likely that I have been evaluating the responses in this biased way, I have reason to think that I should not trust my own evaluations. I should think that it's unlikely that the evidence supports my level of confidence.

In fact, I have evaluated the responses correctly and impartially. I have not been biased in my assessment and I am right to be highly confident that I am popular given my evidence.³³

That the majority of responses I receive from my friends and colleagues are overwhelmingly positive supports my high level of confidence that I am popular. That there's a non-trivial possibility that I have been biased in my evaluation of the responses suggests that my

³²See for instance Brown (1986) and Taylor and Brown (1988).

³³On this example, see Elga (2005).

evidence likely doesn't support my level of confidence. Respecting my evidence seems to require me to be very confident that I'm popular, while also thinking that my level of confidence is too high. This may seem initially problematic. But if by competently evaluating my evidence, propositions like that I am popular or that I am well-liked by the majority of my friends and colleagues become part of my evidence, and my knowledge about the distortion bias can be understood as operating to limit my access to my evidence, then in POPULARITY it can be rational for me to be epistemically akratic. Recognizing that I'm disposed to be biased in a particular kind of way doesn't always require me to revise my levels of confidence, though I should not have the further belief that such a combination of attitudes is permitted on my evidence, given ANTI-AKRASIA*.

Knowledge Iteration. The knowledge iteration principle, KK, states that if you know p , then you know that you know p . Proponents of knowledge iteration often appeal to the infelicity of assertions of the form ' p , but I do not know whether I know p '. If knowledge is the norm of assertion it's hard to see why these assertions should be infelicitous. No contradiction arises when you're in the state $Kp \wedge K\neg KKp$, and in such a state, you should be able to unproblematically make these assertions. Defenders of knowledge iteration suggest that this is evidence for KK, for if knowledge iterates, there will never be cases of $Kp \wedge K\neg KKp$.³⁴ Notice, however, that the epistemic state $Kp \wedge K\neg KKp$ is fragile: You cannot know that you're currently in this state. It's expected that when you're in a fragile epistemic state, asserting the relevant proposition will make salient that you are in such a state – a state which you're forbidden from believing that you are in. The infelicity of asserting ' p , but I do not know whether I know p ' can be explained as arising not from the assertion itself, but rather from the corresponding belief about your epistemic state that tends to follow when you assert propositions of this form.

Substantive and Structural Rationality. Whereas 'substantive rationality' is the kind of rationality typically associated with reasons responsiveness, 'structural rationality' is the kind of rationality typically associated with the relationship between mental states or attitudes. When you fail to believe in accordance with your evidence or when you fail to do what you have most reason to do, you're substantively irrational. When your beliefs are inconsistent or when your preferences are cyclic, you're structurally irrational. Is structural (ir)rationality distinct from substantive (ir)rationality?³⁵

Insofar as EVIDENTIALISM is a genuine requirement of (substantive) rationality and ANTI-AKRASIA is a genuine requirement of (structural) rationality, the conflict between EVIDEN-

³⁴For example Das and Salow (2018, p.3): "... if there are counterexamples to KK, there are fully coherent agents who know that p without being in a position to know that they know this. Plausibly, such agents would be justified (at least sometimes) in judging and asserting that p while refusing to take a stance on whether they know that p . In other words, they would be justified in making the self-undermining or incoherent judgements described above [like 'while it is raining, I'm not willing to take a stance on whether I know that it is'].²⁷ Others, including McHugh (2010), Cohen and Comesaña (2013), and Greco (2015b) appeal to similar considerations in defending KK.

³⁵For discussion see especially Kieseewetter (2017), Lord (2018), and Worsnip (2021).

TIALISM and ANTI-AKRASIA suggests that structural (ir)rationality and substantive (ir)rationality are distinct, for being substantively rational is no guarantee for being structurally rational. But if what rationality requires is ANTI-AKRASIA* and not ANTI-AKRASIA, cases in which you're uncertain what your evidence is or what your evidence supports are not ones in which substantive and structural rationality conflict. That believing in accordance with your evidence can require you to be epistemically akratic is no reason for thinking that structural rationality cannot be reduced to substantive rationality.

7 A Challenge

To end, it's worth reflecting on the initial appeal of ANTI-AKRASIA. If ANTI-AKRASIA isn't a genuine requirement of rationality, why does it appear so compelling? Some speculation: ANTI-AKRASIA appears compelling because it's easy to confuse ANTI-AKRASIA with ANTI-AKRASIA* – it's easy to confuse instances in which you merely believe akratically with instances in which you not only believe akratically, but also believe that your akratic beliefs are supported by your evidence (especially once those beliefs are made salient). This confusion is understandable, as there is some temptation to think that whenever you're rational, your evidence supports believing that you're rational. If this tempting thought is right, ANTI-AKRASIA* would entail ANTI-AKRASIA. But it's not. And on reflection, this is unsurprising. If you think that you should conform your beliefs to your evidence, and you also think that you're not always in a position to determine what your evidence is or what your evidence supports, there is little reason to think that you're always in a position to determine that you are rational whenever you are.³⁶ Given that you ought not believe what's unknowable (or, weaker yet, what's knowably unknowable), something more is true: Sometimes when you're rational, you're required to *avoid* believing that you are. The challenge for advocates of ANTI-AKRASIA is to answer what (if anything) remains uniquely irrational about epistemic akrasia that is left unexplained by ANTI-AKRASIA*.

³⁶Note that my view doesn't place any constraints on your credence about whether you're rational when you're epistemically akratic. In particular, it doesn't require you to have credence 0 that you're rational when you're in a fragile epistemic state. As I see it, this is a feature of the view, for if there are cases in which your evidence supports believing akratically, then presumably you *should* have non-zero credence (perhaps even high credence) that you're in such a case.

References

- [1] Adler, J. 2002. *Belief's Own Ethics*. Cambridge: MIT Press.
- [2] Bird, A. 2007. Justified Judging. *Philosophy and Phenomenological Research*, 74, pp. 81-110.
- [3] Bonjour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- [4] Broome, J. 2013. *Rationality Through Reasoning*. Wiley-Blackwell.
- [5] Brown, J. D. 1986. Evaluations of Self and Others: Self-enhancement Biases in Social Judgments. *Social Cognition*, 4, pp. 353-376.
- [6] Carnap, R. 1950. *Logical Foundations of Probability*. Chicago: Chicago University of Chicago Press.
- [7] Christensen, D. 2007. Epistemology of Disagreement: The Good News. *Philosophical Review*, 116, 187-217.
- [8] Christensen, D. 2010a. Higher-Order Evidence. *Philosophy and Phenomenological Research*, 81, pp. 185-215.
- [9] Christensen, D. 2010b. Rational Reflection. *Philosophical Perspectives*, 24, pp. 121-140.
- [10] Clarke, R. 2013. Belief is Credence One (in Context). *Philosophers' Imprint*, 13, pp. 1-18.
- [11] Clifford, W. K. 1876. The Ethics of Belief. *The Contemporary Review*, 29, pp. 289-309.
- [12] Coates, A. 2012. Rational Epistemic Akrasia. *American Philosophical Quarterly*, 49, pp. 113-124.
- [13] Cohen, S. and J. Comesaña. 2013. Williamson on Gettier Cases and Epistemic Logic. *Inquiry*, 56, 15-29.
- [14] Conee, E. and R. Feldman. 1985. Evidentialism. *Philosophical Studies*, 48, pp. 15-34.
- [15] Das, N. and B. Salow. 2018. Transparency and the KK Principle. *Noûs*, 52, pp. 3-23.
- [16] DeRose, K. 2002. Assertion, Knowledge, and Context. *Philosophical Review*, 111, pp. 167-203.
- [17] Dodd, D. 2017. Belief and Certainty. *Synthese*, 194, pp. 4597-4621.
- [18] Dorst, K. 2019. Higher-Order Uncertainty. In *Higher-Order Evidence: New Essays* eds. M. Skipper and A. Steglich-Petersen. Oxford: Oxford University Press.
- [19] Elga, A. 2005. On Overrating Oneself... and Knowing It. *Philosophical Studies*, 123, pp. 115-124.
- [20] Elga, A. 2007. Reflection and Disagreement. *Noûs*, 41, 478-502.
- [21] Elga, A. 2013. The Puzzle of the Unmarked Clock and the New Rational Reflection Principle. *Philosophical Studies*, 164, pp. 127-139.
- [22] Easwaran, K and B. Fitelson 2015 Accuracy, Coherence, and Evidence. In *Oxford Studies in Epistemology, Volume 5* eds. T. S. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- [23] Feldman, R. 2005. Respecting the Evidence. *Philosophical Perspectives*, 19, pp. 95-119.
- [24] Fogal, D. 2020. Rational Requirements and the Primacy of Pressure. *Mind*, 129, pp. 1033-1070.
- [25] Greco, D. 2014. A Puzzle about Epistemic Akrasia. *Philosophical Studies*, 167, pp. 201-219.
- [26] Greco, D. 2015a. How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives*, 29, pp. 179-201.
- [27] Greco, D. 2015b. Iteration and Fragmentation. *Philosophy and Phenomenological Research* 91, pp. 656-673.

- [28] Greco, D. 2019. Fragmentation and Higher-Order Evidence. In *Higher-Order Evidence: New Essays* eds. M. Skipper and A. Steglich-Petersen. Oxford: Oxford University Press.
- [29] Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- [30] Hawthorne, J., Y. Isaacs, and M. Lasonen-Aarnio. 2021. The Rationality of Epistemic Akrasia. *Philosophical Perspectives*, 35, pp. 206-228.
- [31] Hawthorne, J., D. Rothschild, and L. Spectre. 2016. Belief is Weak. *Philosophical Studies*, 173, pp. 1393-1404.
- [32] Hawthorne, J. and A. Srinivasan. 2013. Disagreement without Transparency: Some Bleak Thoughts. In *The Epistemology of Disagreement: New Essays* eds. J. Lackey and D. Christensen. Oxford: Oxford University Press.
- [33] Hazlett, A. Higher-Order Epistemic Attitudes and Intellectual Humility. *Episteme*, 9, pp. 205-223.
- [34] Hedden, B. and K. Dorst. 2022. (Almost) All Evidence is Higher-Order Evidence. *Analysis*, 82, pp. 417-425.
- [35] Horowitz, S. 2014. Epistemic Akrasia. *Noûs*, 48, pp. 718-744.
- [36] Horowitz, S. and P. Sliwa. 2015. Respecting All the Evidence. *Philosophical Studies*, 172, pp. 2835-2858.
- [37] Huemer, M. 2007. Moore's and the Norm of Belief. In *Themes from G. E. Moore: New Essays in Epistemology and Ethics* eds. Susana Nuccetelli and Gary Seay. Oxford: Oxford University Press.
- [38] Huemer, M. 2011. The Puzzle of Metacoherence. *Philosophy and Phenomenological Research*, 82, pp. 1-21.
- [39] Kelly, T. 2003. Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*, 66, pp. 612-640.
- [40] Kelly, T. 2005. The Epistemic Significance of Disagreement. In *Oxford Studies in Epistemology, Volume 1* eds. T. S. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- [41] Kieseewetter, B. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.
- [42] Kolodny, N. 2005. Why Be Rational? *Mind*, 114, pp. 509-563.
- [43] Kyburg, H. E. 1961. *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- [44] Lasonen-Aarnio, M. 2014. Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research*, 88, pp. 314-345.
- [45] Lasonen-Aarnio, M. 2020. Enkrasia or Evidentialism? Learning to Love Mismatch. *Philosophical Studies*, 177, pp. 597-632.
- [46] Lewis, D. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, 74, pp. 549-567.
- [47] Littlejohn, C. 2018. Stop Making Sense? On a Puzzle about Rationality. *Philosophy and Phenomenological Research*, 96, pp. 257-272.
- [48] Liu, S. 2020. (Un)knowability and Knowledge Iteration. *Analysis*, 80, pp. 474-486.
- [49] Lord, E. 2018. *The Importance of Being Rational*. Oxford: Oxford University Press.
- [50] McHugh, C. 2010. Self-Knowledge and the KK Principle. *Synthese*, 173, pp. 231-257.
- [51] Neta, R. 2019. The Puzzles of Easy Knowledge and of Higher-Order Evidence: A Unified Solution. In *Higher-Order Evidence: New Essays* eds. M. Skipper and A. Steglich-Petersen. Oxford: Oxford University Press.

- [52] Nelkin, D. 2000. The Lottery Paradox, Knowledge, and Rationality. *Philosophical Review*, 109, pp. 373-409.
- [53] Pryor, J. 2018. The Merits of Incoherence. *Analytic Philosophy*, 59, pp. 112-141.
- [54] Radford, C. 1966. Knowledge – By Examples. *Analysis*, 27, pp. 1-11.
- [55] Rinard, S. 2019. Reasoning One's Way Out of Skepticism. In *The Mystery of Skepticism, Volume 2* eds. K. McCain and T. Poston. The Netherlands: Brill.
- [56] Salow, B. 2019. Elusive Externalism. *Mind*, 128, pp. 397-427.
- [57] San, W. K. Disappearing Diamonds: Fitch-Like Results in Bimodal Logic. *Journal of Philosophical Logic*, 48, pp. 1003-1016.
- [58] Scanlon, T. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- [59] Scanlon, T. 2007. Structural Irrationality. In *Common Minds: Themes from the Philosophy of Philip Pettit* eds. G. Brennan, R. Gooden, F. Jackson, and M. Smith. Oxford: Oxford University Press.
- [60] Schechter, J. 2013. Rational Self-Doubt and the Failure of Closure. *Philosophical Studies*, 163, pp. 428-452.
- [61] Schoenfield, M. 2015. Bridging Rationality and Accuracy. *The Journal of Philosophy*, 112, pp. 633-657.
- [62] Shah, N. 2006. A New Argument for Evidentialism. *Philosophical Quarterly*, 56, pp. 481-449.
- [63] Silins, N. 2005. Deception and Evidence. *Philosophical Perspectives*, 19, pp. 375-404.
- [64] Skipper, M. Higher-Order Defeat and the Impossibility of Self-Misleading Evidence. In *Higher-Order Evidence: New Essays* eds. M. Skipper and A. Steglich-Petersen. Oxford: Oxford University Press.
- [65] Smithies, D. 2012. Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research*, 85, pp. 273-300.
- [66] Smithies, D. 2019. *The Epistemic Role of Consciousness*. Oxford: Oxford University Press.
- [67] Sorensen, R. 1988. *Blindspots*. Oxford: Oxford University Press.
- [68] Stalnaker, R. 2006. On Logics of Knowledge and Belief. *Philosophical Studies*, 128, pp. 169-199.
- [69] Sutton, J. 2005. Stick to What You Know. *Noûs*, 39, pp. 359-396.
- [70] Taylor, S. and J. Brown. 1988. Illusion and Well-being: A Social Psychological Perspective on Mental Health. *Psychological Bulletin*, 103, pp. 193-210.
- [71] Titelbaum, M. 2015. Rationality's Fixed Point (or: In Defense of Right Reason). In *Oxford Studies in Epistemology, Volume 5* eds. T. S. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- [72] Unger, P. 1975. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- [73] Weatherson, B. 2019. *Normative Externalism*. Oxford: Oxford University Press.
- [74] Wedgwood, R. 2011. Justified Inference. *Synthese*, 189, pp. 1-23.
- [75] White, R. 2007. Epistemic Subjectivism. *Episteme*, 19, pp. 115-129.
- [76] White, R. 2014. What is My Evidence that Here is a Hand? In *Scepticism and Perceptual Justification* eds. D. Dodd and E. Zardini. Oxford: Oxford University Press.
- [77] Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- [78] Williamson, T. 2014. Very Improbable Knowing. *Erkenntnis*, 79, pp. 971-999.

- [79] Worsnip, A. 2018. The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research*, 96, pp. 3-44.
- [80] Worsnip, A. 2019. Can Your Total Evidence Mislead About Itself? In *Higher-Order Evidence: New Essays* eds. M. Skipper and A. Steglich-Petersen. Oxford: Oxford University Press.
- [81] Worsnip, A. 2021. *Fitting Things Together: Coherence and the Demands of Structural Rationality*. Oxford: Oxford University Press.