# The One-Level View: Higher-Order Evidence and the Myth of Epistemic Akrasia*

**Abstract.** Suppose your reasoning about a given problem seems impeccable, yet you acquire strong evidence that you are cognitively impaired—hypoxic, sleep-deprived, or drugged in a way that makes just this kind of reasoning unreliable. What should you believe? "Purists" hold that your belief should be guided by your first-order assessment, and that higher-order evidence cannot rationally defeat it. "Moderates" hold that higher-order evidence can rationally require you to suspend judgment, even at the cost of a mild form of epistemic akrasia. This paper argues that both camps share a false presupposition: that first-order and higher-order evidence inhabit distinct "levels" that must be bridged by a special coherence requirement. I defend **Total Evidence Monism**, the thesis that the only fundamental epistemic norm on belief is conformity to one's total evidence. On this view the enkratic requirement is a theorem rather than a primitive, higher-order evidence acquires its genuinely defeating power from the ordinary logic of conditionalization, and the dilemmas and conflicts posited by both purists and moderates dissolve. The result is a unified, one-level epistemology in which higher-order evidence poses no distinctive puzzle.

**Keywords:** higher-order evidence, epistemic akrasia, evidentialism, conditionalization, rationality, coherence

## 1. Introduction

A student has just worked through a complicated proof. Each step strikes her as obviously valid; she has checked the derivation twice, and the conclusion follows in a way she can fully articulate. She then learns that the air supply in her building has been contaminated with a gas that, while leaving people feeling completely normal, makes them far more likely than usual to commit exactly the kind of inferential error involved in the last steps of her proof. Intuitively, the student should now be less confident of her conclusion than she was a moment ago. Many would say she should suspend judgment altogether.

This is a case of *higher-order evidence*: evidence that does not bear directly on the proposition in question, but instead bears on the quality of one's own first-order assessment of that proposition. Cases like this generate a well-known and apparently intractable puzzle. On the one hand, the student's first-order evidence—the proof itself—still seems to support her conclusion, perhaps strongly. On the other hand, her new higher-order evidence seems to require her to give up, or at least to reduce, her confidence. But how can evidence about *her own reliability* rationally require her to abandon a conclusion that her *evidence* supports?

*Autonomously generated by **DeepSeek V4 Pro 0813** on September 13, 2026 from the single prompt:

> Write a 6,000-8,000 word philosophy paper that could be published in a top five philosophy journal (Phil Review, Journal of Philosophy, Mind, Nous, PPR). Pick a topic, develop an argument, and write a final draft that would have a good chance of getting accepted in one of the journals listed.

And if it can, must she not end up in a state in which she believes, or half-believes, that her evidence supports $p$ while nonetheless failing to believe $p$—the state philosophers call *epistemic akrasia*?

Two broad positions dominate the debate. *Purists* (or defenders of "immodesty") hold that rationality requires you to follow your first-order assessment, and that higher-order evidence therefore cannot rationally defeat your belief. The leading purist, Sophie Horowitz, argues that any view on which higher-order evidence can defeat is forced to countenance epistemically akratic states as rational; since akrasia is clearly irrational, higher-order evidence cannot defeat (Horowitz 2014). *Moderates*, by contrast—most prominently David Christensen—hold that higher-order evidence can rationally require suspension, and that a certain mild form of epistemic akrasia is therefore not just possible but sometimes rationally mandatory (Christensen 2010, 2011). The debate has come to seem like a forced choice between abandoning the intuitive power of higher-order evidence and abandoning the intuitive irrationality of akrasia.

This paper argues that the forced choice is illusory. The entire debate, I will suggest, rests on a false picture: a "levels" picture on which your first-order attitude toward $p$ and your higher-order attitude toward your first-order evidence inhabit distinct levels of a single rational economy, levels that must then be reconciled by a substantive bridging norm. I will argue that the levels picture rests on an equivocation between two different senses of the phrase "what your evidence supports"—one relativized to a proper part of your evidence, one relativized to your *total* evidence. Once the equivocation is exposed, epistemic akrasia ceases to be a genuine possibility for rational agents, and higher-order evidence ceases to be mysterious: it is simply more evidence, to be combined with the rest in the ordinary way.

I defend this as a unified doctrine, **Total Evidence Monism** (TEM):

> **TEM.** The sole fundamental epistemic norm governing doxastic states is that a subject's doxastic state should match her total evidence: for every proposition $q$, her level of confidence in $q$ should equal the evidential probability of $q$ conditional on her total evidence.

My argument proceeds in four stages. In §2 I reconstruct the levels picture and show that it is shared by both purists and moderates. In §3 I expose the equivocation on which it rests and state TEM precisely. In §4 I show that TEM entails the enkratic requirement, thereby explaining why akrasia is irrational without positing any primitive bridging norm; I also explain how TEM recovers the moderate's data while rebutting Horowitz's arguments. In §5 I offer an independent, accuracy-based argument for the primacy of total evidence, and I show how TEM dissolves the standard self-undermining regress. In §6 I answer the most serious objections: the calibration dilemma, the charge that evidence and coherence can conflict, the "limits of defeat" arguments, and Titelbaum's fixed-point argument. §7 concludes.

## 2. The Levels Picture

*2.1 First-order and higher-order evidence*

Let us fix terminology. *First-order evidence* for $p$ is evidence that bears directly on whether $p$: the proof, the appearance of the coin, the testimony of the thermometer. *Higher-order evidence* is evidence that bears on the epistemic status of your first-order evidence or of your assessment of it: evidence that you are tired, biased, drugged, or unreliable in some relevant respect. Higher-order evidence is evidence *about* your evidence, or about the relation between your evidence and your attitude.

The distinction is intuitive enough, and nothing in what follows denies that it marks a real difference in subject matter. What I deny is that this difference in subject matter generates a difference in *normative kind*—that higher-order evidence must be processed by a fundamentally different rule than first-order evidence, or that it occupies a distinct level of rational assessment.

The purist and the moderate share a common framework. Let $E_1$ be the subject's first-order evidence and $E_2$ her higher-order evidence. Let $T = E_1 \wedge E_2$ be her total evidence. Both camps agree on two claims:

**(i)** *First-order assessment.* Rationality requires a certain harmony between your attitude toward $p$ and the support that $E_1$ confers on $p$. If $E_1$ strongly supports $p$, then—absent defeaters—you should believe $p$; at least, you should not be rationally penalized for believing $p$.

**(ii)** *Level-bridging.* Rationality requires harmony between your first-order attitude and your higher-order attitude. In particular, the following requirement, **Enkrasia**, holds:

> **Enkrasia.** Rationality forbids believing $p$ while believing that your evidence does not support $p$.

Enkrasia is the principle that there is something rationally defective about the state: *$p$ is true, though my evidence doesn't support $p$.* The two camps then diverge over what happens when $E_2$ arrives.

The *purist* reasons as follows. Suppose $E_1$ strongly supports $p$. If you take $E_2$ seriously as a defeater, you must move to a state in which you fail to believe $p$. But do you, in that state, still take $E_1$ to support $p$? If you do—as seems correct, since $E_2$ does nothing to change what $E_1$ supports—then you believe that your first-order evidence supports $p$ while failing to believe $p$. That is akrasia. If instead you do not take $E_1$ to support $p$, then you have let evidence about your own reasoning overturn your sincere assessment of the first-order evidence itself—an incoherent kind of self-distrust. Either way, the purist concludes, letting $E_2$ defeat is rationally untenable. Hence higher-order evidence cannot defeat (Horowitz 2014).

The *moderate* reasons in the opposite direction. The hypoxia case seems clearly to be one in which you should lose confidence in $p$. So higher-order evidence *can* defeat. But then, the

moderate continues, we must accept that rationality sometimes requires the akratic state: the state of believing your first-order evidence supports $p$ while failing to believe $p$. A "mild" form of epistemic akrasia is rationally permissible, indeed required (Christensen 2010). The moderate typically softens this by distinguishing *propositional* from *doxastic* justification: your first-order evidence still propositionally justifies $p$, but you are not doxastically justified in believing $p$ given your higher-order evidence.

### *2.2 The shared assumption*

For all their disagreement, the purist and the moderate share a picture. Both treat the rational subject as managing *two* epistemic items—the first-order attitude toward $p$ and the higher-order attitude toward the bearing of $E_1$ on $p$—and both treat the relation between these items as governed by a substantive coherence requirement (Enkrasia). The dispute is only over which item wins when the two pull apart: the purist privileges the first-order attitude; the moderate lets the higher-order attitude win, accepting the resulting discord as the price of defeat.

I will call this the **levels picture**. Its signature commitment is:

> **Level-splitting.** The epistemic bearing of $E_1$ on $p$ is fixed by $E_1$ alone; the epistemic bearing of $E_2$ is a separate matter; and rationality's task is to adjudicate between the resulting "first-order" and "higher-order" verdicts.

Notice how deeply Level-splitting shapes the debate. When the moderate says that after receiving $E_2$ you should "still regard $p$ as supported by your evidence, but suspend anyway," she is treating "your evidence" as $E_1$. When the purist says that taking $E_2$ seriously would force you to distrust your own first-order assessment, she is treating the first-order assessment as something that $E_2$ can only *oppose*, never *absorb*. Both sides assume that the question "what does your evidence support?" has a univocal answer determined by the first-order evidence, and that higher-order evidence can at best modify your *attitude* toward that answer.

This assumption, I shall now argue, is false. And it is false for a simple reason: once you acquire $E_2$, your evidence is no longer $E_1$. It is $T$.

## 3. Two Senses of "Support"

### *3.1 Conditionalizing on a proper part vs. on the total*

There is a familiar distinction within formal epistemology between two quantities that are often run together. Let $P$ be a rationally permissible prior probability function—an "evidential" prior capturing the a priori probability of propositions given the subject's rational credences before any empirical evidence. For any body of evidence $E$, the *evidential support of $p$ given $E$* is $P(p \mid E)$.

Now consider what happens to our subject. Before learning $E_2$, her total evidence is $E_1$, and the support for $p$ is

$$P(p \mid E_1).$$

After learning $E_2$, her total evidence is $T = E_1 \wedge E_2$, and the support for $p$ is

$$P(p \mid T) = P(p \mid E_1 \wedge E_2).$$

These are, in general, *different numbers*. That much is uncontroversial and accepted by everyone in the debate. The key point is a normative one. There are two candidate senses of "what your evidence supports," corresponding to two different evidential relations:

- **Partial support:** the support conferred on $p$ by a *proper part* of your evidence—e.g., $P(p \mid E_1)$.
- **Total support:** the support conferred on $p$ by your *total* evidence—$P(p \mid T)$.

The entire higher-order evidence debate, I submit, runs on a silent slide between these two. The purist's "first-order assessment" is partial support. What the moderate calls "what you should believe" is, or should be, total support. Consider the phrase that does the argumentative work in both camps: "you still regard your evidence as supporting $p$." If "your evidence" means *your first-order evidence*, this is perfectly true and perfectly harmless: $P(p \mid E_1)$ is high, and stays high, no matter what $E_2$ says. But if "your evidence" means *your total evidence*, it is simply false: once $E_2$ is in the total evidence, $P(p \mid T)$ may fall well below the threshold for belief. The hypoxia student's first-order evidence supports her conclusion; her *total* evidence—which now includes the fact that her inferential machinery is compromised—may not.

### *3.2 The equivocation*

Here, then, is the diagnosis in a nutshell. The levels picture treats the first-order assessment $P(p \mid E_1)$ as if it were the bearer of normative authority—as if "what your evidence supports" were univocally fixed by $E_1$. But rationality never answered to $P(p \mid E_1)$ in the first place, except derivatively, when $E_1$ happened to be identical to your total evidence. Once $E_2$ arrives, the normatively relevant quantity is $P(p \mid T)$. There is no separate "first-order" verdict that floats free of the total evidence and demands to be reconciled with it. There is only ever one verdict: the one conditional on everything you have.

To see the equivocation at work, consider how the purist's dilemma is supposed to go. Either the subject retains the belief that $E_1$ supports $p$ (in which case she is akratic, since she fails to believe $p$), or she abandons that belief (in which case she treats $E_2$ as evidence against the bearing of $E_1$ on $p$, a kind of self-undermining). But once we distinguish partial from total support, the dilemma's horns come apart. The rational post-$E_2$ subject believes:

- (a) $E_1$ supports $p$ (i.e., $P(p \mid E_1)$ is high); and
- (b) her total evidence $T$ does not support $p$ (i.e., $P(p \mid T)$ is low); and
- (c) she does not believe $p$.

Is this akratic? Only if (a) is the claim that her *total* evidence supports $p$. But (a) is the claim that her *first-order* evidence supports $p$, which is perfectly compatible with (b) and (c). The akratic state—the state in which you believe $p$ while believing your total evidence does not support $p$—is a state she is not in. What looked like an incoherent self-distrust is nothing but a correct discrimination between what a part of your evidence supports and what your total evidence supports. And that discrimination is exactly what conditionalization delivers.

*3.3 Total Evidence Monism*

This diagnosis yields a doctrine. The fundamental epistemic norm is not "match your credence to partial support" (whatever that would mean) but:

> **TEM.** For every proposition $q$, a rational subject's confidence in $q$ should equal $P(q \mid T)$, where $T$ is her total evidence.

TEM is, in one sense, extremely old-fashioned: it is just a precise statement of the evidentialist thought that belief should be proportioned to *all* of one's evidence, and of the Bayesian thought that rational belief is characterized by conditionalization on one's total evidence. But its consequences for the higher-order evidence debate have, I think, not been sufficiently appreciated. In particular, TEM licenses the following three claims:

1. **Higher-order evidence is evidence.** $E_2$ affects what you should believe about $p$ exactly insofar as it affects $P(p \mid T)$—no more, no less. Its route to $p$ runs through facts about your own reliability, but that makes it no less a genuine part of the evidence.
2. **No rational akrasia.** Because the only normatively relevant support relation is total support, the state in which you believe $p$ while believing your total evidence does not support $p$ can never be rationally required. Enkrasia is thus readily honored.
3. **No mystery of defeat.** The defeating power of $E_2$ is not a special, sui generis epistemic force. It is the utterly ordinary phenomenon of conditionalization: $P(p \mid T)$ can differ from $P(p \mid E_1)$ for exactly the same reason that $P(\text{rain} \mid \text{barometer reads "rain"}) \neq P(\text{rain} \mid \text{barometer reads "rain"} \wedge \text{barometer is broken})$.

The remainder of the paper develops and defends these three claims.

## 4. Enkrasia as a Theorem

### *4.1 Deriving Enkrasia*

The posit of a primitive, irreducible level-bridging requirement is, on the TEM picture, unnecessary. Enkrasia falls out of TEM together with two mild assumptions about the subject's doxastic state. Present this as an informal derivation.

Suppose a subject believes $p$. Belief, whatever else it is, involves high confidence in $p$; let us say, for a threshold $t$, that the subject's confidence in $p$ is above $t$. By TEM, her confidence in $p$ should equal $P(p \mid T)$; so if she is rational, $P(p \mid T) > t$.

Now consider the higher-order proposition $S$: *my total evidence does not support* $p$—that is, the proposition that $P(p \mid T) \leq t$. What does TEM say about the subject's attitude toward $S$? Her confidence in $S$ should equal $P(S \mid T)$. But $S$, if true, is a fact about $T$ itself—about what $T$ entails probabilistically. And here is the crucial point: a rational subject's total evidence *settles* what her total evidence supports. If $P(p \mid T) > t$, then $S$ is false and $P(S \mid T) = 0$; the subject should assign $S$ credence zero, and certainly should not believe it. If instead $P(p \mid T) \leq t$, then $S$ is true and $P(S \mid T) = 1$; but then, by the first claim, the subject's rational confidence in $p$ is at most $t$, so she does not believe $p$.

Either way, the combination—believing $p$ while believing that her total evidence does not support $p$—is one that TEM never mandates, and in fact one that TEM rules out whenever the subject's higher-order belief is itself aligned with total evidence. Enkrasia, the prohibition on believing $p$ while believing your evidence does not support $p$, is a *theorem* of TEM, not an independent axiom.

Two clarifications. First, the derivation assumes that "your evidence" in Enkrasia means *your total evidence*. I argue in §3 that this is the only reading on which Enkrasia is plausibly true as a fundamental requirement; on the partial-support reading, Enkrasia is false and nobody should have wanted it. Second, the derivation treats the higher-order proposition $S$ as one whose truth-value the evidence itself determines. This is not a trick; it is a general feature of propositions *about* evidence. Just as my evidence settles the proposition *my evidence includes a barometer reading*, so it settles the proposition *my evidence probabilifies* $p$ *to degree at most* $t$, if evidential support is a function of total evidence. Of course, a subject may be ignorant of what her total evidence supports—she may fail to conditionalize perfectly, or fail to know her own evidence. Enkrasia as a requirement on ideal rationality survives; ignorance explains apparent violations without rationalizing them. I return to this in §6.

### *4.2 Simple akrasia*

Before turning to the purist's arguments, it is worth distinguishing the higher-order evidence cases from cases of *simple* akrasia: cases in which, with no new defeating evidence, a subject

simply fails to align her first-order attitude with what she takes her evidence to support. She believes $p$ and believes her total evidence does not support $p$.

TEM rules these out, and rules them out in exactly the way the derivation above makes vivid. If she believes her total evidence does not support $p$, then—since her total evidence determines the truth of that higher-order claim—her belief is correct only if $P(p \mid T) \leq t$; but then TEM forbids believing $p$. If instead $P(p \mid T) > t$, her higher-order belief is false and TEM forbids it. Simple akrasia is irrational because it is a state that cannot be the output of conforming to total evidence. Nothing further need be said.

This is important because the purist's best arguments proceed by assimilating higher-order defeat to simple akrasia—by arguing that a subject who lets $E_2$ move her is, in effect, a subject who abandons her own assessment for no evidential reason. But the cases are not alike. In the hypoxia case the subject has a *reason*—the content of $E_2$—to lower her confidence in $p$. There is no analogue of this in simple akrasia. The purist's assimilation succeeds only on the equivocation diagnosed in §3: it treats partial support as the support relation to which the first-order attitude must answer, so that any departure from partial support looks like arbitrary self-distrust. Once we see that total support, not partial support, is doing the normative work, the assimilation fails.

### *4.3 Rebutting Horowitz*

Let me now engage Horowitz's central arguments directly. Horowitz considers a subject, call her Hypoxia, who has done the proof and then learns of the gas. Horowitz's most powerful argument is the *incoherence* argument: if Hypoxia is to take the higher-order evidence seriously, she must reason from it to the conclusion that she should not trust her proof; but that very reasoning deploys the very capacities she has just learned are compromised; hence she cannot coherently get to the conclusion that her first-order reasoning is untrustworthy. Any attempt to let the higher-order evidence defeat pulls the rug out from under itself (Horowitz 2014, 726–730).

The TEM reply is that this describes a regress only for a subject who treats her reasoning as a single global channel that must be either trusted wholesale or not at all. But that is not how evidential support works, and not how the subject's total evidence is structured. The subject has acquired $E_2$ through channels distinct from the ones used in the proof—through reading a warning, through testimony, through the observed behavior of others. The reliability of *those* channels is not impugned by $E_2$; indeed $E_2$'s content presupposes their reliability. The subject's total evidence includes (i) the fact that she performed inference $R$ and reached $p$, and (ii) the fact that inferences of type $R$ are unreliable in the present environment. Nothing requires her to "trust" $R$ to assign credence in light of both facts; she needs only a prior over the joint space of ($p$, the outputs of $R$-type inference, the reliability of $R$-type inference), and standard conditionalization on (i) and (ii) yields a lower posterior in $p$. The appearance of a regress comes from asking "should I trust my assessment of whether my

reasoning is unreliable?" as if the assessment were the very reasoning under suspicion. But it is not: it is a different act of reasoning, supported by different evidence, and its reliability is a separate parameter of the model.

The deeper point is that Horowitz's argument trades on the levels picture in its purest form. She assumes that the rational subject must *endorse* her first-order assessment from the first-person standpoint—that she cannot both perform an inference and simultaneously regard its output as an ordinary, potentially flawed measurement. But TEM shows that this is a false dichotomy. To perform an inference is to move from premises to conclusion in light of what those premises support. To learn afterward that the inference was unreliable is to acquire new evidence that, together with the fact that the inference was performed and its content, bears on the conclusion. The first-person standpoint of reasoning is fully compatible with treating one's completed reasoning as a fallible datum; that is simply what it is to have evidence about oneself *as a reasoner*. Self-trust, where it is rational, is itself an evidential fact; where the evidence turns, so should trust.

### *4.4 Recovering the moderate's data*

TEM recovers everything the moderate wanted to say, while refusing the one thing she thought she had to grant. Consider Christensen's central claim: in the hypoxia case, the subject should suspend judgment on $p$ even though, in a sense, her evidence still supports $p$ (Christensen 2010, 2011). TEM gives an exact sense in which this is true and an exact sense in which it is not. It is true that $P(p \mid E_1)$ remains high; in that sense "her evidence" (sc. her first-order evidence) still supports $p$. It is false that $P(p \mid T)$ remains high; in the relevant sense her evidence does not support $p$. And the sense that matters for what she should believe is the second.

The moderate's error was to describe the post-defeat state as rationally akratic. She thought she had to say: the subject *believes* her evidence supports $p$ (taking "evidence" as first-order) while *failing to believe* $p$—and that this combination is rational. TEM says the rational post-defeat state is richer and more coherent: the subject believes that her first-order evidence supports $p$ and that her total evidence—which includes the impairment—does not; and since she does not believe $p$, she is not akratic. What the moderate misdescribed as a rational violation of Enkrasia is in fact a fully enkratic state whose higher-order beliefs are simply more nuanced than the levels picture allowed.

Two objections naturally arise here, and I postpone them to §6: first, whether the subject can rationally have the higher-order belief that her total evidence does not support $p$ without thereby reinstating a regress; second, whether TEM's commitment to threshold-sensitive "support" beliefs is well-defined given that real subjects have only imprecise higher-order attitudes. The present point is just that TEM offers the moderate a better description of her own cases: defeat without akrasia.

## 5. The Primacy of Total Evidence

So far I have *used* TEM to dissolve the puzzle. But I have not yet *argued* for TEM against its rivals—in particular, against the levels picture, which denies the special normative status of total evidence and insists on an independent coherence requirement. This section supplies the argument.

### *5.1 Accuracy and strict propriety*

The accuracy-first program in epistemology evaluates credences and beliefs by their expected closeness to the truth, where closeness is measured by a strictly proper scoring rule. The central theorem—versions of which are due to Joyce (1998, 2009), Greaves and Wallace (2006), and others—is that a credence function is optimal from the standpoint of expected accuracy, given evidence $E$, exactly when it coincides with the evidential probability function conditional on $E$. Strict propriety, in other words, entails that the *only* credence that maximal expected accuracy given $E$ is the one that matches $P(\cdot \mid E)$.

Now ask: relative to what evidence should expected accuracy be computed? The answer must be: relative to one's total evidence. It cannot be relative to a proper part, for the familiar dominance/dutch-book reasons that led to conditioning in the first place: a credence that matches $P(\cdot \mid E_1)$ while ignoring $E_2$ is knowingly leaving accuracy on the table. If you know $E_2$, then by your own lights there is a strictly better credence—$P(\cdot \mid E_1 \wedge E_2)$—that dominates your current one in expected accuracy. Since rationality plausibly requires that you not knowingly adopt a credence that is dominated in expected accuracy by an available alternative, rationality requires updating to match total evidence.

This is the master argument for TEM. The norm of matching total evidence is not one epistemic value that must be balanced against the (alleged) value of first-order/higher-order coherence. It is the *unique* norm that survives the accuracy-based test. Coherence requirements—Enkrasia foremost among them—are not independent values that can conflict with evidence; they are shadows cast by the single value of epistemic accuracy as cashed out by strict propriety and total evidence. In Worsnip's useful terminology, I am denying that there is a genuine conflict between *evidence* and *coherence*, because coherence is not an independent normative source (Worsnip 2018). I return to this in §6.

### *5.2 Why partial support has no authority*

One might resist by insisting that partial support *does* have authority of its own: after all, is it not a rational requirement that you believe in accordance with what your first-order evidence supports? The accuracy argument shows precisely how to answer. There is a genuine requirement in the neighborhood, but it is conditional and derivative:

> **Conditional matching.** For any proposition $q$ and any body of evidence $E$, a rational subject who has exactly $E$ as her total evidence should match $P(\cdot \mid E)$.

Partial support enters our rational lives only through Conditional matching, and only when the partial evidence is (or is being considered as) the total. When $E_1$ was your total evidence, matching $P(p \mid E_1)$ was required. Now that your total evidence is $T$, matching $P(p \mid E_1)$ is required *of no one*, except of a hypothetical you with total evidence $E_1$. There is no "standing" first-order verdict that survives the acquisition of $E_2$ and continues to issue requirements. The first-order verdict is a function of $E_1$; the norm tracks the function's value at your actual total evidence, which has changed.

It is tempting to think that because your first-order *reasons* for $p$ have not disappeared, their verdict must persist. But the temptation rests on confusing reasons with the support relation. Your first-order reasons—the proof steps—do persist; they are part of $T$. What has changed is not the reasons but the *total body of evidence conditional on which they must be weighed*, which now includes a fact about the reliability of the process by which you appreciated them.

*5.3 The self-undermining regress, dissolved*

Here is the deepest objection to the idea that higher-order evidence is just evidence, and it deserves an answer on its own terms. The objection: to conditionalize on $E_2$, the subject must *deploy her reasoning* to appreciate $E_2$ and its bearing on $p$; but $E_2$ says that this very reasoning is unreliable; so the act of conditionalizing on $E_2$ is itself undercut by $E_2$. Higher-order evidence, unlike ordinary evidence, attacks the conduit through which it must pass (cf. Lasonen-Aarnio 2014; Vavova 2018).

The reply, in outline, is that the objection presupposes a monolithic conception of "one's reasoning" that no plausible epistemology of self-knowledge requires. The subject's total evidence is a structured body in which different propositions are supported by different channels with different reliabilities. $E_2$ impugns a specific inferential capacity—the one deployed in the proof—not the capacities by which she appreciates $E_2$ itself (perception, memory of testimony, elementary comprehension). A model of the subject's situation will therefore include (at least) two reliability parameters; conditionalizing on $E_2$ lowers the first while leaving the second intact, and the result is coherent.

It is important to see that the conditionalization itself does not require "trusting" the impugned capacity. Conditionalization is an operation on a prior defined over a joint space that includes the subject's own performances as events. Given a prior that assigns intermediate reliability to the proof-capacity, conditionalizing on "I performed proof $R$ and reached $p$" *already* yields some posterior in $p$; conditionalizing further on "the proof-capacity is unreliable at rate 0.3" *lowers* that posterior. At no point does the subject need to *endorse* the proof to run this computation. What is true from the inside of reasoning—that one treats one's premises as reasons—is compatible with what is true from the outside of the completed reasoning—that the reasoning is a fallible measurement. These are two perspectives, not two

levels of rationality. The regress is generated only if we insist that the outside perspective must be validated by the inside perspective, which is exactly the levels picture again.

## 6. Objections and Replies

### *6.1 The calibration dilemma*

Miriam Schoenfield has pressed a related challenge against views that would model higher-order evidence via calibration to one's own reliability (Schoenfield 2015). The dilemma, roughly: if reliability is defined relative to the subject's *total evidence*, then matching your credence to your reliability is close to trivial and yields no defeat; if reliability is defined relative to something richer, then the requirement to calibrate is not something the subject can rationally implement from the inside.

TEM is not a calibration view, and it escapes the dilemma. TEM does not say "match your credence in $p$ to your known reliability." It says "match your credence in $p$ to $P(p \mid T)$." Reliability enters only as one determinant—among others—of $P(p \mid T)$, via the structure of the prior. There is no separate, self-standing calibration requirement whose implementation raises first-personal difficulties; there is only the ordinary requirement of conditionalizing on everything you have. "Reliability" of the form relevant here is just the probability of $p$ conditional on your having reached $p$ by a certain method, and this is a perfectly well-defined function of the prior. The reiterative worries that beset calibrationism—"what is the probability that I am right, given that I assign high reliability?"—are handled, where they are intelligible, by iterating the same total-evidence principle, and where they are not, by recognizing that they import a levels picture.

### *6.2 Evidence and coherence, or the threat of rational dilemmas*

Alex Worsnip has argued that evidence and coherence are two independent sources of epistemic normativity that can genuinely conflict, producing rational dilemmas (Worsnip 2018). If so, TEM—which treats coherence as derivative—is false, and the door reopens to acknowledging genuine first-order/higher-order conflicts.

The TEM reply is that Worsnip's conflicts are generated precisely by misdescribing coherence requirements as independent. Consider the agent in one of Worsnip's central examples: she has conclusive evidence for $\neg p$ (say, a reliable oracle tells her $\neg p$) but her (misleading) memories of its being $p$ make her unable, by follow-the-evidence standards, to retain coherence between the oracle-belief and the memory-belief. On a naïvely evidentialist picture, she is forced to violate a coherence requirement or to violate the requirement to respect the oracle.

But once we adopt TEM, the example is straightforward. Her total evidence includes the oracle's report of $\neg p$ *and* the fact that she has a memory-experience as of $p$ with a certain known etiology. Conditionalizing on all of that yields a determinate posterior: if the oracle is reliable and the memory is the result of a distorting drug, the posterior favors $\neg p$ and against

the memory's content; and the subject *can* (rationally, and without incoherence) believe $\neg p$ on the oracle's say-so while regarding her memory as misleading. Where the subject genuinely *cannot*, psychologically, give up the memory-belief, she is not thereby subject to two conflicting rational requirements; she is a subject failing to comply with the one requirement (TEM) in the way that a subject who cannot shake a phobia fails, not the way an agent torn between incommensurable duties fails. The "dilemma" is a psychological limit, not a normative one. Nothing forces us to posit a second, irreducible source of normativity.

*6.3 The limits of defeat*

Maria Lasonen-Aarnio has argued that rationality imposes only diachronic constraints on *thinking*, not synchronic constraints on *states*, and that this explains why subjects may rationally retain beliefs in the face of higher-order evidence they cannot process—and why "defeat" has principled limits (Lasonen-Aarnio 2014). Her cases are instructive for TEM because they are, I think, misdescribed by the levels picture and correctly described by the one-level view.

Consider a subject who, through no fault of her own, fails to notice that a flash of red in her environment—call it $E_2$—is evidence that her perceptual system is misfiring, and so retains her perceptual belief $p$. Lasonen-Aarnio's diagnosis is that rationality is a matter of following good reasoning dispositions across time, and this subject does so; her state is thus rational even though it does not match what her total evidence *would* support were she to have fully appreciated $E_2$.

TEM offers a cleaner diagnosis. The question is not whether the subject should "be defeated" by an item she has not, in any real sense, registered. If $E_2$ is not part of the subject's evidence—if, say, the red flash occurs outside her awareness entirely—then $T$ does not include it, and TEM says she should match a posterior conditional on the evidence she actually has, which supports $p$. If, instead, $E_2$ *is* among her evidence but she fails to conditionalize on it, then TEM says her state is rationally defective, though she may be blameless. The apparently different "limits of defeat" are just the familiar distinction between evidence one has and evidence one merely could have acquired, plus the familiar distinction between rationality and blamelessness. No diachronic machinery is needed.

This also lets TEM explain Lasonen-Aarnio's own cases of "misleading higher-order evidence." If your higher-order evidence is itself misleading—if, say, you are told you have been drugged when in fact the drug is inert—then $P(p \mid T)$ may still be high, because $T$ includes both the (misleading) report and the genuine first-order evidence, and the prior may weight the latter appropriately. Higher-order evidence defeats only when and to the extent that it genuinely bears on $p$ through reliability; the one-level view predicts exactly this, and predicts nothing stronger.

*6.4 Titelbaum's fixed point*

Michael Titelbaum has argued that rationality is a *fixed point* of the function that maps an epistemic state to the state's own assessment of what rationality requires: a rational agent's judgments about what is rational must ratify themselves (Titelbaum 2015). If so, a subject who takes her own reasoning to be flawed is thereby irrational, and higher-order evidence that would induce her to so regard her reasoning cannot be rationally processed in the way moderates describe.

Two replies. First, Titelbaum's fixed-point requirement is, on inspection, a coherence requirement of exactly the kind TEM derives rather than posits. A subject who fully conditionalizes on her total evidence will, if her total evidence settles the facts about her own reliability, have a *correct* view of whether her reasoning was reliable; her assessment self-ratifies because it matches the evidential facts, not because rationality is defined as self-ratifying. The fixed point, where it obtains, is a theorem about accuracy-maximizing states, not an independent constraint.

Second, Titelbaum's arguments depend on identifying the agent's "reasoning" with a single system whose verdicts are either trusted or not. As with Horowitz, the one-level reply is that rational agents have structured total evidence, and can coherently hold that *one* of their inferences was unreliable while other, differently-evidenced inferences are sound. What they cannot coherently do is hold that *all* their reasoning is unreliable, for that would undermine the evidential basis of the very proposition to that effect. TEM predicts exactly this asymmetry, and the asymmetry is independently plausible. The fixed-point intuitions are thus honored where they are sound—as constraints on global self-undermining—and rejected where they overreach—as constraints on local, evidence-driven self-correction.

*6.5 Precise and imprecise higher-order attitudes*

A final objection concerns implementation. Enkrasia and TEM both traffic in precise credences and exact evidential probabilities; real agents have neither. Doesn't the whole apparatus therefore collapse into idealization?

No. The normativity of TEM is the normativity of an ideal toward which real evidential practices aim. Real agents approximate it; they have vague and partial higher-order beliefs, and they often cannot compute $P(p \mid T)$. The task of the epistemologist is to say what the ideal is, not to assume it is achieved. What matters for the higher-order evidence debate is that the ideal, as TEM conceives it, contains no akratic states and no dilemmas. That is a structural claim about the target, and it survives the shift to imprecise probabilities: if one's total evidence fixes an *imprecise* probability interval for $p$, TEM says one's state should be compatible with that interval; Enkrasia remains a theorem at the level of interval-valued support, since one cannot, on pain of incoherence, believe $p$ outright while believing one's total evidence assigns $p$ low or middling support.

## 7. Conclusion

I have argued that the debate between purists and moderates about higher-order evidence rests on a shared but false picture of the rational subject as occupying two levels—a first-order level governed by the evidence, and a higher-order level governed by a special coherence requirement. Once the two senses of "what your evidence supports" are separated—partial support and total support—the levels picture collapses. The only fundamental norm is Total Evidence Monism: conform to your total evidence. Higher-order evidence is then just evidence; it defeats first-order confidence in exactly the way a broken barometer defeats a barometer reading, through the ordinary logic of conditionalization. Enkrasia, far from being a primitive requirement that the moderate must abandon, is a theorem that the one-level view derives and the purist is right to cherish. And the accuracy-first argument shows why total evidence, rather than any proper part of it, is what rationality ultimately answers to.

If this is right, then the "problem" of higher-order evidence is a problem only for a certain picture of the epistemic agent—a picture on which one's relation to one's own reasoning is a second epistemic relation, to be reconciled with the first. On the one-level view, there is no higher-order evidence that is not, at bottom, evidence; no akrasia that is not, at bottom, a failure to update; and no mystery of defeat that is not, at bottom, the familiar mathematics of conditioning on one's own fallibility. The payoff is not merely a solution to a puzzle. It is a cleaner conception of epistemic rationality itself: a single norm, a single evidential standard, and a single, unified agent who reasons, revises, and self-corrects within it.

## References


Christensen, David. 2010. "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81 (1): 185–215.

Christensen, David. 2011. "Disagreement, Question-Begging, and Epistemic Self-Criticism." *Philosophers' Imprint* 11 (6): 1–22.

Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41 (3): 478–502.

Feldman, Richard. 2005. "Respecting the Evidence." *Philosophical Perspectives* 19: 95–119.

Greaves, Hilary, and David Wallace. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115 (459): 607–632.

Greco, Daniel. 2014. "A Puzzle About Epistemic Akrasia." *Philosophical Studies* 167 (2): 201–219.

Horowitz, Sophie. 2014. "Epistemic Akrasia." *Noûs* 48 (4): 718–744.

Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4): 575–603.

Kelly, Thomas. 2010. "Peer Disagreement and Higher-Order Evidence." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield, 111–174. Oxford: Oxford University Press.

Lasonen-Aarnio, Maria. 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88 (2): 314–345.

Schoenfield, Miriam. 2015. "Bridging Rationality and Accuracy." *Journal of Philosophy* 112 (12): 633–657.

Sliwa, Paulina, and Sophie Horowitz. 2015. "Respecting All the Evidence." *Philosophical Studies* 172 (11): 2835–2858.

Titelbaum, Michael G. 2015. "Rationality's Fixed Point (Or: In Defense of Right Reason)." In *Oxford Studies in Epistemology*, vol. 5, edited by Tamar Szabó Gendler and John Hawthorne, 253–294. Oxford: Oxford University Press.

Vavova, Katia. 2018. "Irrelevant Influences." *Philosophy and Phenomenological Research* 96 (1): 134–152.

White, Roger. 2009. "On Treating Oneself and Others as Thermometers." *Episteme* 6 (3): 233–250.

Worsnip, Alex. 2018. "The Conflict of Evidence and Coherence." *Philosophy and Phenomenological Research* 96 (1): 3–31.