

# Has Horowitz Split Level-Splitting?

SAVANNAH LEON\*

*University of California, Berkeley*

**Abstract** What is to be done when first- and higher-order evidence point in opposite directions concerning the truth about  $p$ ? The traditional response goes that ideally rational agents ought to privilege one evidential order over the other, such that an agent's belief that  $p$  co-varies with her total evidence. But the level-splitter zigs where others zag. Since each evidential order appears perfectly good in isolation, she supposes her credences should be partitioned accordingly. On penalty of believing against her total evidence, she responds to the pull of both evidential orders. In other words, she is epistemically akratic. Sophie Horowitz has recently argued that level-splitting views are almost universally irrational. To show as much, she points to some cases of peer disagreement where a pro-akrasia verdict requires (irrationally) concluding that  $S$ 's evidence is misleading. The purpose of this paper is to deny that an on-off conception of agent-specific defeaters is called for: that is, I argue that peer disagreement need not necessarily banish first-order evidence to the realm of the misleading, and that a different approach is available to the pro-akrasia crowd.

## 1 Introduction

In 'Epistemic Akrasia,' Sophie Horowitz argues that while rational epistemic akrasia can be warranted in special cases, a pro-akrasia solution in standard cases is too intuitively costly to be right (Horowitz 2014). The problem of akratic belief states has inspired a triad of potential solutions in the literature of epistemology of disagreement: two traditional, one contemporary. Traditional responses reject the notion of rational epistemic akrasia. That is, they reject the view that one ought to believe  $p$  while simultaneously believing that  $p$  is unsupported by evidence. A contemporary response

---

\*Savannah Leon is now a fourth year philosophy student at the University of California, Berkeley. Her philosophical interests include epistemology, philosophy of language, and some formal topics. Besides doing philosophy, she also enjoys travelling, reading poetry, and cooking. She plans on applying to philosophy programs in the UK and US as a prospective graduate student after finishing at Berkeley.

is that such belief states are only *prima facie* problematic and may even be rationally required.

The paper is structured as follows. In §1, I introduce Horowitz's Sleepy Detective Problem and three responses to it: the traditional (anti-akrasia) *conciliatory* and *steadfast* verdicts and the newer (pro-akrasia) *level-splitting* verdict. In §2, I motivate rational epistemic akrasia. I foreshadow Horowitz's reasons for thinking that level-splitting is irrational in most cases, yet rationally required in others. In §3, I discuss two further cases and give an overview of Horowitz's account of evidential uncertainty and how evidence can be considered either truth- or falsity-guiding. In §4, I give her argument in support of level-splitting's nearly universal irrationality, which features cases of peer disagreement where pro-akrasia requires (irrationally) concluding that one's evidence is misleading. In §5, I argue that these cases hinge on an assumption that a certain species of defeater commits level-splitters to forming the belief that their evidence is misleading, and that since this assumption is false, it doesn't follow that level-splitters must conclude that their evidence is falsity-guiding. In doing this, I indirectly offer an alternative explanation for why akratic belief states may sometimes be rational despite the worry posed by Horowitz.

## 2 What is epistemic akrasia?

'Akrasia' has classically meant *weakness of will*, such as in cases where S acts against her better judgment.<sup>1</sup> Correspondingly, 'epistemic akrasia' refers to cases wherein S arguably believes *p* against her better judgment. (For example, when her belief that *p* appears inadequately supported by her available evidence.) While not all-encompassing (and purposely somewhat imprecise, given current debates as to just what epistemic akrasia is), this definition shall serve as a reasonable point of departure in understanding what it is to be epistemically akratic.

The paper will proceed, as Horowitz does, with the evidentialist approach to understanding epistemic akrasia. Provisionally, I will consider the puzzle of whether epistemically akratic belief states are to be thought rational as one which can be solved or dissolved by settling on how an agent should apportion her belief that *p* with respect to her total evidence concerning *p*. If we suppose that two or more crucial parts of an agent's available evidence is both in favor of and against believing that *p* (and that she is unable to suspend judgment), and the agent is ultimately uncertain as to whether *p*, it otherwise remains unclear how and in which circumstances her belief that *p* could

---

1. An attempt to account for the 'paradoxical irrationality' of akratic agents is of course discussed in the Protagoras, and also in Davidson (1982). Another account, that of Levy (2018), argues that what we take to be (epistemically) akratic states are an agent's mistaken belief that they believe that *p*. Levy differentiates between belief that *p* and a(n) agent's indistinct first-order 'beliefy' representation(s) that *p*.

be rationally required.

Part of this puzzle's traction is owed to emerging controversies about evidence. The controversy currently at hand turns on a tension between two sorts of evidence: first- and higher-order evidence. The first term is maybe the most familiar: 'first-order evidence' is that which bears directly on  $p$ 's truth value. Standard sources of first-order evidence include (non-exhaustively) perception and memory. Contrastingly, 'higher-order evidence' is typically thought to be evidence *about* one's evidence, or evidence 'bearing on the functioning of one's rational faculties, or on the significance of other evidence that one has' (see Horowitz, forthcoming). Higher-order evidence, then, can speak to how  $S$  ought to interpret her first-order evidence, and often crops up in the form of testimony (as in cases of peer disagreement) or even reflection upon one's own epistemic state (such the recognition of impairment or lack of expertise). For instance, higher-order evidence's impact on total evidence (especially concerning defeat, i.e., whether higher-order evidence can undercut or rebut first-order evidence) remains at large in recent literature (Christensen 2010; Feldman 2009; Lasonen-Aarnio 2014), and what differentiates higher-order evidence from first-order evidence is also an open question.

There are many cases that serve as excellent candidates for framing the kind of evidential tension which might rationalize epistemic akrasia.<sup>2</sup> For simplicity, I'll use Horowitz's: its key details bear centrally and specifically on her criticisms of level-splitting and the subject of this paper. The case:

**Horowitz's Sleepy Detective** Sam is a police detective, working to identify a jewel thief. He knows he has good evidence—out of the many suspects, it will strongly support one of them. Late one night, after hours of cracking codes and scrutinizing photographs and letters, he finally comes to the conclusion that the thief was Lucy. Sam is quite confident that his evidence points to Lucy's guilt, and he is quite confident that Lucy committed the crime. In fact, he has accommodated his evidence correctly, and his beliefs are justified. He calls his partner, Alex. 'I've gone through all the evidence,' Sam says, 'and it all points to one person! I've found the thief!' But Alex is unimpressed. She replies: 'I can tell you've been up all night working on this. Nine times out of the last ten, your late-night reasoning has been quite sloppy. You're always very confident that you've found the culprit, but you're almost always wrong about what the evidence supports. So your evidence probably doesn't support Lucy in this case.' Though Sam hadn't attended to his track

---

2. Consider, for example, Alvin Plantinga's 'letter filching case' (1986), where a man accused of stealing a letter has (arguably excellent) first-order evidence that he didn't take it—in this case, his memory of walking in the woods. He is in fact correct about his woods memory. As with the Sleepy Detective case, there's nevertheless a large body of higher-order evidence against him: he's done similar things before, and an extremely reliable person testifies that she witnessed the theft.

record before, he rationally trusts Alex and believes that she is right—that he is usually wrong about what the evidence supports on occasions similar to this one (Horowitz 2014, 2).

To specify: Sam’s first-order evidence (hereafter ‘FOE’) is the aforementioned codes, letters, and photographs. Sam’s higher-order evidence (hereafter ‘HOE’) consists in Alex’s testimony that Sam is probably unable to properly interpret FOE. Taking together Sam’s first- and higher-order evidence provides a (rough) picture of his ‘total evidence.’<sup>3</sup> Next I will spell out what each possible verdict has to say about how these evidential orders should interact with one another, if at all.

One might think that the two evidential orders *don’t* (or shouldn’t) interact, and furthermore that Sam might be rational in believing both that Lucy is the culprit while also accepting Alex’s testimony that he is often wrong in situations such as these. This option permits rational epistemic akrasia.<sup>4</sup> The view that akratic attitudes can be rationally required in cases like Sleepy Detective is what Horowitz dubs the ‘level-splitting’ position. For our purposes, the shortened ‘pro-akrasia’ will often be used. If the detective is epistemically akratic, then he’ll continue to believe that Lucy is the jewel thief while believing that his total evidence doesn’t support this. On pro-akrasia, then, he should split his levels of confidence and hold onto the first-order belief that  $p$  given FOE and the higher-order belief about his unreliability given HOE.

The traditional ‘anti-akrasia’ response is to deny that epistemically akratic states are rational. On this view, it is often thought that epistemic levels should *never* operate separately. An anti-akrasia proponent would apply this notion in the Sleepy Detective case by requiring that Sam must base his belief in Lucy’s guilt on his total evidence. For Sam, this would mean that he must either steadfastly remain confident that Lucy is guilty on the basis of FOE, or else be persuaded to reduce his confidence in her guilt given Alex’s testimony that he ought to doubt his initial conclusion to the contrary.

### 3 Level-splitting

Horowitz calls the pro-akrasia view that epistemic levels should operate separately ‘level-splitting’ (Horowitz 2014), which is to be properly differentiated from epistemic akrasia. Rather than merely labelling cases of ‘divergence between first- and higher-

---

3. It’s not uncontroversial that Sam’s total evidence might contain more or less than what is said here; this can vary depending on one’s views concerning which evidence an agent ought to consider.

4. I’m grateful to the aforementioned anonymous referee for pointing out that although level-splitters hold the view that evidential orders should operate separately and that this goes some way in rationalizing epistemic akrasia, one’s commitment to the interaction of evidential orders isn’t simultaneously a commitment to level-splitting in epistemic akrasia cases. One might think, for example, that epistemic akrasia is irrational without this entailing that evidential orders shouldn’t influence one another, and vice-versa.

order attitudes' (Horowitz 2014), level-splitting views rationally require such divergences.<sup>5</sup>

In other words: level-splitting is a normative position, whereas epistemic akrasia is used descriptively. I'd also like to note that Horowitz perceives that first- and higher-order belief mismatches as coming in degrees:

- **MODERATE LEVEL-SPLITTING** prescribes being highly confident that  $p$  despite having high confidence that your evidence that doesn't support your degree of confidence in  $p$ .
- **EXTREME LEVEL-SPLITTING** recommends high confidence in  $p$  while also being rationally highly confident that (a) your evidence doesn't support  $p$ , (b) your evidence supports low confidence in  $p$ , or (c) your evidence supports  $\neg p$ .

These distinctions are a sticking point in section 4, where I'll zero in on the belief state depicted in 'C'. Unless otherwise specified, it can be assumed for now that when I refer to level-splitting, I have in mind the gamut of akratic states given above.

A defense of pro-akrasia is founded in the thought that if either evidential order seems perfectly good to us in isolation, our belief state ought to reflect this somehow. On the basis of both FOE + HOE, the level-splitter believes that  $p$  and believes that there's something fishy about her total evidence for  $p$ . Again, many grounds for fishiness exist: misleading evidence, poorly-interpreted evidence, insignificant evidence, etc.

Horowitz argues that, barring a complex and much-discussed case, epistemic akrasia is universally irrational.<sup>6</sup> This caveat leads her to distinguish between two kinds of cases, **STANDARD** and **NONSTANDARD**. We shall feature an example or two from each in sections 2.3 and 2.4. The cases are classed according to a couple of contrasting features: (1) different types of uncertainty, and (2) opposite background expectations about how our evidence should point to the truth about  $p$ .

**Uncertainty** In standard cases, a would-be akratic agent is uncertain about what her total evidence supports. The difficulty here lies in discerning which order of evidence is really getting it right about  $p$ . In nonstandard cases, even if she can be sure of what her evidence should support, the problem is that she can't be sure of what her evidence *is*.

**Truth- and falsity-guiding evidence** Standardly, and nearly unanimously, we expect our evidence to be **TRUTH-GUIDING**, so that 'when it justifies high confidence in a

5. The two might easily be seen as interchangeable. I wish to avoid this confusion.

6. The exceptional case is that of Williamson's 'irritatingly austere' clock (from Williamson 2011, 2014 discussed at length in Elga 2013); see p. 6 for Horowitz's adaptation.

proposition, that proposition is usually true, and when it justifies low confidence in a proposition, that proposition is usually false' (Horowitz 2014). But in unusual cases, we can have the background expectation that our evidence—whatever it is—will vary *falsely* with the proposition it is meant to support. That is, we expect that it will be **FALSITY-GUIDING**: it'll support high confidence in a false proposition, and low confidence in a proposition we think is probably true. Horowitz concludes that special cases like these are rare, but plausibly justify pro-akrasia.

Horowitz considers a number of epistemologists who argue that epistemic akrasia is rationally required.<sup>7</sup> We shall consider a pair of cases from two authors in particular: Weatherson (n.d.), who uses a familiar case of justified moral akrasia, and Williamson (2011, 2014), whose case is purely epistemic. The abridged version of each:

**Weatherson's Kantian Professor** By way of sophisticated and persuasive argumentation, suppose your Kantian professor has given you good evidence to believe that lying is categorically wrong. Nevertheless, when a murderer inquires as to your roommate's whereabouts, you lie, since lying is what you ought to do. (Horowitz 2014)

**Williamson's Long Deduction** Suppose a rational agent comes to know a long series of claims and deduces their conjunction, *C*. She's done so competently, but she realizes that since oftentimes memory and logical ability are limited, people in her situation often make inferential errors while completing long deductions. It's then highly probable on her evidence that she herself has made such an error, and thus that she doesn't know the conjunction. Still, given that she's competently deduced *C*, she knows *C*: its evidential probability is 1. It's nevertheless highly probable on her evidence about fallibility during long deductions that she doesn't know *C*. So, she should be highly confident in *C* despite her high confidence that she doesn't know *C*. (Horowitz 2014)

Borrowing a line from Lewis (1996), it seems a level-splitting agent can rationally 'properly ignore' evidence across epistemic levels: that is, form a kind of provisional belief that *p*. Acknowledging the evidential force of FOE + HOE rids us of the drawback of ignoring good evidence—in these cases, a properly performed proof and a well-formulated normative claim.

---

7. Whatever the authors themselves may make of level-splitting, I'll follow Horowitz in proceeding as though the views presented here commit them to it.

## 4 Problems for level-splitting

Horowitz gives more than a few examples where making the choice to split epistemic levels goes terribly wrong. Central to this paper, however, is one case which is supposed to demonstrate that Sam—were he epistemically akratic—would have strange beliefs indeed about where his evidence points. In particular, Horowitz says: ‘If [the detective] takes both “Lucy is the jewel thief” and “my evidence doesn’t support Lucy” as premises, it seems he can engage in some patently bad reasoning’ (Horowitz 2014).

Horowitz’s argument is as follows. Suppose that Sam trusts Alex’s testimony and forms the belief that the odds are 1:9 that Lucy is guilty. Suppose further that despite this, he remains confident that she is the culprit. Horowitz reasons that the detective must then think that, given such low odds, he ‘just got lucky’ about his true belief. Given that he rationally trusts Alex, Sam *should* be confident FOE doesn’t support Lucy’s guilt. (That is, he should have low confidence that  $p$ .) His HOE therefore pushes him towards high confidence that  $\neg p$ . A plausible (extreme) pro-akrasia reading of his *total* evidence, then, is that it’s falsity-guiding: that is, it supports high confidence in a false proposition,  $\neg p$ .

Horowitz gives an analogous case, *SPIKED COFFEE*, where Sam has received HOE (Alex’s testimony) that his coffee’s been spiked with a reason-distorting serum (Horowitz 2014). His coffee has not in fact been spiked. Although he has the reasonable background expectation that his FOE would be truth-guiding under normal conditions, his partner’s testimony causes him to form the belief that he can’t properly evaluate his evidence. Nevertheless, he has properly evaluated his evidence, and  $p$  is a true belief. As in the above example, however, he also believes that his evidence is misleading on the basis of HOE. This, argues Horowitz, *shouldn’t* be enough to cause Sam to form the belief that his FOE is misleading, so it’s irrational to believe as much given HOE.

In what follows, I’ll examine the claim that extreme pro-akrasia verdicts of the form ‘my evidence doesn’t support  $p$ ’, or ‘my evidence supports low confidence in  $p$ ’ must land level-splitters in the predicament given above. The Spiked Coffee case, when paired with its ‘just got lucky’ predecessor, is supposed to point out the following absurdity. It would be absurd for an agent to form the belief that her evidence supports  $\neg p$ —that is, that her evidence is misleading—merely on the basis of a defeater bearing on her agent’s diminished capacities (a ‘self-doubting defeater’).

## 5 Spiked Coffee, revisited

Horowitz writes:

If Level-Splitting is right, and extreme cases of epistemic akrasia can be rational in Sleepy Detective, there is nothing wrong with Sam's concluding that his evidence is misleading in this way. But there *is* something wrong with Sam's concluding that his evidence is misleading in this case. This suggests that there is something wrong with Level-Splitting. (Horowitz 2014).

For the record, I'm confident Horowitz is right on the order of it being silly for the detective, spiked coffee or no, to draw the conclusion that his evidence is falsity-guiding. It doesn't seem at all plausible that a claim bearing on an agent's capacities should tarnish first-order evidence to the effect that it supports  $\neg p$ .

I'll argue against Horowitz that a level-splitter needn't conclude from a self-doubting defeater ('your coffee's been spiked, so you should doubt your ability to interpret your evidence!') that her first-order evidence is misleading. This would, as Horowitz has claimed, be irrational. I leave open the possibility that that some level-splitter or other might draw such a conclusion, but by my lights, there is an available alternative.

Might we say instead that the detective's confidence that Lucy is guilty should be *reduced* by his knowledge of his track record? Alternatively, we could just as well leave open the possibility that the detective might refrain from believing anything about whether his first-order evidence points to falsehood or truth. It's not clear why self-doubting defeaters must have an on-off effect on belief in such a way as to pressure extreme level-splitters to immediately interpret their evidence as only truth-guiding or falsity-guiding, but not something in between. The defeater might be classified as undercutting, or alternatively as neither raising nor lowering the probability that  $p$ . These interpretations sidestep the need for self-doubting defeaters to outright rebut the proposition that one's first-order evidence supports  $p$ .

If, for example, I were to espouse an extreme level-splitting view and then interpret my evidence under the influence of Irish coffee, I wouldn't form the belief that my FOE is—or must be—falsity-guiding. Depending on the amount of Irish coffee involved, I might reduce my confidence as to where my evidence points, or alternatively take the 'wait-and-see' approach until morning. Still, I needn't believe that my spiked coffee points away from the truth about  $p$ , and neither does the detective. He could refrain from believing anything about whether his first-order evidence points towards truth or falsehood, or merely reduce his confidence that Lucy is guilty upon becoming aware of his track-record.

In short, level-splitters aren't committed to the view that self-doubting defeaters are necessarily falsity-guiding, or that such defeaters should always have the effect of transforming one's evidence that  $p$  into evidence that  $\neg p$ . When paired with other compelling strands of evidence, a creeping suspicion of one's own unreliability *might* lead an agent to wonder about what her first-order evidence really supports, but then again it may not.

The same point evinces, I think, in another stock example of epistemic akrasia. If we consider Williamson's Long Deduction case, we might also think that a tendency to make inferential errors of this sort shouldn't serve the same evidential role as an added line in a proof that points away from your conclusion that *C*. Still, the extremely akratic logician doesn't appear committed to concluding that her evidence is misleading. This is because she can just as easily maintain high confidence that *C* while believing that her evidence could be something unnervingly short of truth-guiding.

If my analysis is correct, the detective needn't worry about his partner's testimony showing his FOE to be falsity-guiding; i.e., that it perplexingly supports Lucy's innocence ( $\neg p$ )<sup>8</sup>, or somehow shrouds the truth about *p* in a thinner sense than rebutting it.<sup>9</sup> In much the same way, it would also helpfully tie up the loose end that the extreme level-splitter's answer in long deduction (as initially stated) isn't doomed to the same fate.

## 6 Conclusion

On Horowitz's view, an immediate problem with level-splitting is that it permits irrationally concluding that one's evidence is misleading in cases like Sleepy Detective. This, she argues, is problematic: clearly the Sleepy Detective Problem's would-be akratic agent can avoid being misled and can even point to a belief of his that should be revised given his total evidence. I've concluded that if level-splitting is correct, then it's not the case that this evidence *must* be permissibly interpreted as misleading. There may be internecine disagreements among level-splitters as to why and whether this kind of move can be rational. In this case, level-splitters need not *all* think it permissible to interpret evidence in this fashion, so Horowitz's criticism doesn't seem to count against the entire position.

The Sleepy Detective Problem's very setup permits believing (rationally) that one's first-order evidence might not be misleading. Why? I suspect this is because the self-doubting defeater contained in Alex's testimony is a long way off from serving as positive evidence for the devastating conclusion that Sam, were he to continue believing *p*, has been knowingly misled and yet still believes *p*. By my lights, this also tells a plausible story about why a level-splitter's belief state can be stable in cases like Williamson's Long Deduction. Indeed, it's hard to see how just about any close reading of the case would lead one to believe that the akratic logician's higher-order evidence supports  $\neg p$ . For again, if I competently completed a long proof, the likelihood of my

---

8. While I find this perplexing given how I've read the evidential support relation as being 'HOE rationalizes low confidence in *p*' (and so on, and so on), I realize that the detective could very well have it that HOE makes Lucy's guilt less likely than her innocence. Thanks to an anonymous commentator at UCSD for pressing me on this.

9. Thanks to Jennifer Carr for a rich discussion of this point and for an elegant formulation of it.

getting it wrong about the conclusion wouldn't be enough to induce credence  $x$  that something *other* than  $C$  holds: I'd just have good higher-order evidence to suspend belief that  $C$ , or reduce my credence in the proposition 'I know  $C$ .' I'd hardly be forced to revise my belief state such that I have low confidence in  $C$  being the answer.

Even if Horowitz has split level-splitting, I've disagreed with the portion of her strategy wherein an on-off conception of defeaters is called for: that is, if peer disagreement *must* banish first-order evidence to the realm of the misleading. I've offered an alternative that can better explain how akratic belief states might be rational. This point has perhaps a small yield in terms of the broader debate about responding to counterevidence, but if correct, it extricates split-friendly epistemologists from the view that their first-order evidence must be misleading due to higher-order evidence to the contrary. It allows the level-splitter to retain confidence that her evidence is truth-guiding and vindicates the evidence they've evaluated from being doomed to falsity-guiding status in paradigm cases.<sup>10</sup>

## References

- Alston, William P. 1980. 'Level-Confusions in Epistemology'. *Midwest Studies in Philosophy* 5 (1): 135–50.
- Carr, Jennifer. n.d. 'Imprecise Evidence Without Imprecise Credences'. Unpublished.
- Christensen, David. 2010. 'Higher-Order Evidence'. *Philosophy and Phenomenological Research* 81 (1): 185–215.
- Davidson, Donald. 1982. 'Paradoxes of Irrationality'. In *Problems of Rationality*, edited by Donald Davidson. Oxford University Press.
- Egan, Andy, and Adam Elga. 2005. 'I Can't Believe I'm Stupid'. *Philosophical Perspectives* 19 (1): 77–93.
- Elga, Adam. 2013. 'The Puzzle of the Unmarked Clock and the New Rational Reflection Principle'. *Philosophical Studies* 164 (1): 127–39.
- Feldman, Richard. 2009. 'Evidentialism, Higher-Order Evidence, and Disagreement'. *Episteme* 6 (3): 294–312.
- Horowitz, Sophie. 2014. 'Epistemic Akrasia'. *Noûs* 48 (4): 718–44.

---

10. I'm indebted to Jennifer Carr and the audience at Columbia University. Milo Phillips-Brown, Kevin Dorst, and Dee Payton have provided extensive comments on this paper for which I am also grateful. Enormous thanks are in order to an anonymous referee at St Andrews whose constructive feedback led to many positive changes in the paper.

- Horowitz, Sophie. Forthcoming. 'Predictably Misleading Evidence'. In *Higher-Order Evidence: New Essays*, edited by Mattias Skipper and Asbjørn Steglich-Petersen. Oxford University Press.
- Lasonen-Aarnio, Maria. 2014. 'Higher-Order Evidence and the Limits of Defeat'. *Philosophy and Phenomenological Research* 88 (2): 314–45.
- Levy, Neil. 2018. 'You Meta Believe It'. *European Journal of Philosophy* 26 (2): 814–26.
- Lewis, David. 1996. 'Elusive Knowledge'. *Australasian Journal of Philosophy* 74 (4): 549–67.
- Plantinga, Alvin. 1986. 'The Foundations of Theism: A Reply'. *Faith and Philosophy* 3 (3): 313–96.
- Weatherson, Brian. n.d. 'Do Judgments Screen Evidence?' Unpublished.
- Williamson, Timothy. 2011. 'Improbable Knowing'. In *Evidentialism and its Discontents*, edited by T. Dougherty, 147–64. Oxford University Press.
- . 2014. 'Very Improbable Knowing'. *Erkenntnis* 79 (5): 971–99.