

Fragmentation and Higher-Order Evidence

Penultimate Draft, forthcoming in *Higher-Order Evidence: New Essays*

Daniel Greco

The concept of higher-order evidence—roughly, evidence about what our evidence supports—promises epistemological riches; it has struck many philosophers as necessary for explaining how to rationally respond to disagreement in particular, and to evidence of our own fallibility more generally. But it also threatens paradox. Once we allow higher-order evidence to do non-trivial work—in particular, once we allow that people can be rationally ignorant of what their evidence supports—we seem to be committed to a host of puzzling or even absurd consequences. My aim in this paper will be to have my cake and eat it too; I’ll present an independently motivated framework that, I’ll argue, lets us mimic the particular case judgments of those who explain how to accommodate evidence of our fallibility by appeal to higher-order evidence, but without commitment to the absurd consequences.

1 Road Map

My strategy will be as follows. I’ll start by introducing the idea of higher-order evidence (HOE), along with some of the examples which it is often thought to illuminate. In particular, my focus will be on the claim that, due to limited higher-order evidence, we are sometimes rationally uncertain about what our evidence supports. I’ll then review

For helpful comments and discussion, thanks to two anonymous referees for this volume.

some puzzles this claim engenders. I'll divide them in two categories. First, there are the synchronic puzzles. In particular, it seems to entail that epistemic akrasia—believing some claim, while simultaneously believing that you shouldn't believe it—can be rational. But epistemic akrasia seems paradigmatically irrational (Greco, 2014b; Horowitz, 2014; Littlejohn, 2018). Second, there are a host of diachronic puzzles. Having presented an overview of reasons to be skeptical of the idea that we can be rationally uncertain about what our evidence supports, I'll introduce an alternative strategy for making sense of the cases. In a nutshell, it involves the idea that evidence can be available for some purposes or tasks, but not others, and that we often do better to explicitly relativize talk about a subject's evidence to those purposes or tasks which the evidence is available to guide.¹ I'll then argue that when this independently motivated framework is applied to the sorts of cases typically discussed in the higher-order evidence literature, it lets us offer plausible descriptions of cases that are usually characterized in terms of rational uncertainty of what our evidence supports, without the puzzles that characterization brings in its wake.

2 Higher-Order Evidence

As a first pass, we might say that HOE is evidence about evidence.² But this rough characterization obscures crucial distinctions—*some* species of HOE, so defined, raise distinctive epistemological problems that others do not. For instance, if I read in a newspaper that paleontologists have discovered powerful evidence that the Cretaceous extinction was not, after all, caused by a meteor impact, I've plausibly gotten higher-order evidence in this minimal sense—I've gotten evidence (what I read in the newspaper) about evidence (what the paleontologists have discovered). But, for reasons that should

¹This is a special case of the idea that our mental lives are “fragmented”. For some defenses of various versions of this view, see Lewis (1982), Stalnaker (1984, 1991, 1999), Egan (2008), Gendler (2008), Rayo (2013, Chapter 4.), Greco (2014a,b), and Elga and Rayo (Manuscript A, B).

²For instance, Richard Feldman (2009) characterizes HOE as “evidence about the existence, merits, or significance of a body of evidence.”

Fragmentation and Higher-Order Evidence

become clear, this sort of case doesn't raise the tricky issues that have typically been the focus of the recent literature on HOE. For that reason, I'll work in this essay with a narrower characterization of HOE. In the target sense, HOE for a subject *S* is evidence that bears on what *S*'s evidence supports. To receive higher-order evidence, in the sense I'm interested in, is for her to receive:

1. Evidence about which body of evidence *S* herself has, or
2. Evidence about evidential support relations—in particular, evidence about which propositions are supported to which degrees by a body of evidence that, for all *S* knows, is her own.

In this narrower sense, when I read about paleontological discoveries, I don't seem to get any higher-order evidence. I get new evidence to be sure, but that doesn't seem helpfully characterized as either (a) evidence about what my evidence is, or (b) evidence about what that evidence supports.³ While this characterization of HOE may look oddly disjunctive, I don't think it is. 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 I might know which body of evidence is mine, I might fail to know what that given body of evidence supports.⁴ The two-clause characterization of HOE mirrors the two potential sources of uncertainty about what one's evidence supports.⁵

Ultimately—counterintuitively, I think—I'll suggest that epistemologists can and should avoid appeal to HOE, as characterized above. That is, so long as we're careful

³See also Christensen (2010a), who offers a different characterization of HOE, but which would agree with mine in not counting the newspaper case. In particular, he claims that what's distinctive about HOE is that it “rationalizes a change of belief precisely because it indicates that my previous beliefs were rationally sub-par.” (p. 185) And that's not going on in the newspaper case—in coming to suspend judgment about whether a meteor strike caused the cretaceous extinction after reading the newspaper story, I needn't believe that my previous confidence on that point was rationally sub-par.

⁴Of course, I might also be ignorant in both ways.

⁵So I agree with Dorst (Forthcoming), who frames his discussion in terms of higher-order uncertainty, rather than higher-order evidence.

not to equivocate on “evidence”, we can safely theorize in a framework in which a subject’s evidence always settles (a) what her evidence is, and (b) what that evidence supports, and so in which HOE has no non-trivial role to play. But for reasons that should become clear, it’s extremely tempting to equivocate on “evidence”. Appeals to HOE in the recent literature should typically be understood as involving a kind of subtle equivocation—a subject may lack evidence (in one sense) about what her evidence (in some other sense) is, or supports. But I’m getting ahead of myself—reasons to accept these surprising claims will only emerge later in the paper. For now, having offered a general characterization of HOE, I’ll turn to some familiar examples in which it seems natural and fruitful to theorize in terms of it.

2.1 Examples

2.1.1 Disagreement

Probably the best-known examples that have been thought to illustrate the importance of HOE concern disagreement. Here’s a typical one:⁶

JUROR: You’re a juror in a complex civil case in which a great deal of evidence was presented. On the basis of that evidence, you believe that the accused was at fault, so your tentative plan is to find for the plaintiff. And let’s stipulate that your reading of the evidence is, in fact, maximally rational; if an ideal intellect—an epistemic angel—were to look at your evidence, she’d come to the same conclusion you did. But when you sit down to deliberate, you find that many of the other jurors think otherwise—they believe that the accused was not at fault. Upon

⁶Perhaps the most famous example of disagreement is David Christensen (2007b) “Restaurant Case”, but I’m avoiding it on purpose. In short, that case raises the problem of logical omniscience, in addition to problems concerning disagreement. While I do think the diagnosis I’ll ultimately offer of this case could extend to that one, it would be complicated, and would go via the view defended in Elga and Rayo (Manuscript A, B.) Because I don’t have the space here to summarize their treatment of the problem of logical omniscience, nor how I would appeal to it in treating cases of disagreement, I’ll stick to cases of disagreement that don’t pose that problem.

Fragmentation and Higher-Order Evidence

further discussion, you realize it's not that anybody is bringing to bear unshared, background evidence—rather, you just disagree about what conclusion the shared body of evidence presented at court points to. Moreover, the other jurors seem like generally sensible people, and they're not committing any obvious blunders or misunderstandings. At this point, what should you think about whether the accused was at fault?

How might we analyze this case—in particular, capturing plausible verdicts about what it's rational to believe both before and after hearing about the other jurors' views—in terms of HOE? Let's suppose that you are rational in believing that the accused is at fault prior to hearing about the other jurors' views, and then rational in suspending judgment on that question afterwards—this pattern of judgments is typical of “conciliationism”, in the lingo.⁷ For what it's worth, this will make my task harder—while there are extant views that downplay the significance of HOE, they tend to do so by treating other jurors' views as irrelevant to what, in fact, your evidence supports concerning the liability of the accused.⁸

In this case it's natural to say that you and the other jurors know what the evidence is. We can suppose you all have a very good memory, and that if you were asked about whether the body of evidence presented at court included this or that fact, you would always be able to respond accurately. And none of that changes when the other jurors express their views. So HOE in the first sense—evidence about which body of evidence is yours—doesn't seem to be at issue. But, *prima facie*, it's plausible that HOE in the second sense *is* at issue. While you may know what your evidence is, because you're not an ideal evaluator of evidence—or at least, if you are, you don't know *that*—it's not transparent to you what the evidence supports concerning fault. As a matter of fact you judged correctly that it supported the belief that the accused was at fault, but you hold

⁷E.g., Christensen (2007b), Elga (2007), Feldman (2007).

⁸I have in mind Titelbaum (2015) and Smithies (2015), both of whom hold that any body of evidence maximally supports all *a priori* truths about evidential support relations.

that belief about evidential support only defeasibly, and rightly so. When you learn that the other jurors came up with a different answer, your belief is defeated. You obtained misleading higher-order evidence to the effect that your evidence doesn't support the claim that the accused was at fault after all. And if you're rational, you'll respond by suspending judgment, both about fault, and about what your evidence supports concerning fault.

The fact that a reliable inquirer with the same evidence as you disagrees with you suggests you may have made a mistake in judging what your evidence supports. But disagreement isn't the *only* way you can get evidence of such a mistake. The literature abounds with cases in which, due to learning that you've been drugged,⁹ or are suffering from oxygen deprivation,¹⁰ or sleep deprivation,¹¹ or some other malady, you get evidence that is naturally interpreted as suggesting that you made a mistake concerning evidential support relations—you took a body of evidence to support some proposition that, in fact, it does not—even when you have in fact made no such mistake.

2.1.2 Limited Discrimination

While I characterized HOE as coming in two species, aimed at rectifying two sorts of ignorance—ignorance about what your evidence is, and ignorance about evidential support relations—the cases discussed in the previous subsection only concerned the latter species. What about the former? They tend to turn on facts about our limited perceptual discriminatory capacities. Here's an example from Salow (2017b), who is adapting a case from Williamson (2011). I use Salow's version because I'll later appeal to his explanation of what's unattractive about the HOE-based interpretation of the case:

Imagine that you are faced with an unmarked clock, with a single hand

⁹Christensen (2007a).

¹⁰Elga (Ms.)

¹¹Christensen (2010a).

Fragmentation and Higher-Order Evidence

that can point in any one of 60 slightly different directions. Your ability to discriminate where it is pointing is good, but not unlimited. If you are to be reliable in your judgements, you need to leave yourself a margin of error. For example, if the hand is in fact pointing at 53 minutes, you can reliably judge that it is pointing somewhere between 52 and 54 (inclusive), but are unreliable about claims stronger than that. The same is true of every other position the hand could be in.

It is somewhat natural to identify your evidence with the strongest claim about the hand's position which you can reliably get right...If the hand is in fact pointing at 53, my evidence will be that it is within [52, 54]; and if it is pointing at 52, my evidence will be that it is within [51, 53].

What's crucial about the setup, for present purposes, is that you're not in a position to know what your evidence is.¹² Suppose, e.g., the clock is in fact pointing at 53, so your evidence is that it is between 52 and 54 (inclusive). We may assume that you know the general facts about the setup—you know about your impressive but not unlimited discriminatory capacities, know that the hand could be pointing anywhere from 1 to 60, etc. But in that case, if you were to know that your evidence is that the hand is between 52 and 54, then you would be in a position to know that the hand is actually pointing at 53, contradicting the stipulations of the case—by stipulation, that is your evidence only when the clock is pointing precisely at 53. So you must not know that your evidence is that the hand is between 52 and 54. Rather, if the strongest thing you know about where the hand is pointing is that it's between 52 and 54, then instead of knowing exactly what your evidence is, there are three possibilities compatible with your knowledge concerning what your evidence is—either your evidence is that the hand is between 51 and 53, or your evidence is that it's between 52 and 54, or your evidence is

¹²Like both Salow and Williamson, I'll assume $E = K$ —your evidence is what you know. So failing to know what your evidence is amounts to lacking evidence concerning what your evidence is.

that it's between 53 and 55.

3 Puzzles

Now that we've been introduced to the notion of HOE, and to the sorts of examples it's typically used to illustrate, it's time to turn to the difficulties it brings in its wake.

3.1 Epistemic Akrasia

Perhaps the most frequently discussed puzzle in this neighborhood concerns “epistemic akrasia”—believing some claim while also believing that one shouldn't believe it. This is often generalized to involve other sorts of mismatch between one's first-order doxastic attitudes, and one's attitudes concerning which attitudes one should have.¹³ Epistemic akrasia can seem paradigmatically irrational. But as many writers have pointed out, if we accept the evidentialist claim that an agent should believe whatever her evidence supports, and we also accept the existence of non-trivial higher-order evidence—that is, if we accept that agents can be rationally uncertain either about what their evidence is, or what that evidence supports—then it's hard to avoid holding that agents should sometimes be epistemically akratic.¹⁴ The connection has been spelled out at greater length elsewhere, so I'll just give the flavor of the idea here.

First, suppose agents can be rationally uncertain about what their evidence is. Say you don't know whether your total evidence is A, B, or C. You rationally think that each is equally likely to be your evidence. A supports the belief that P, but B and C do not. In fact, your evidence is A. If you have all the beliefs that are supported by your evidence, you'll believe that P (after all, that's supported by A, which is your evidence), while also believing that you probably shouldn't believe that P (after all, you think

¹³Worsnip (2015) and Rinard (Forthcoming) offer similar generalizations of the anti-akrasia constraint, both in terms of having some doxastic attitude while failing to believe that the attitude is supported by one's evidence.

¹⁴See Salow (2017a, Appendix B), Dorst (Ms.).

it's more likely than not that your evidence is either B or C, neither of which support believing that P).¹⁵

Things go much the same if agents can be rationally uncertain about what their evidence supports. Suppose you know what your total evidence is. But there are three hypotheses about what that evidence supports—A, B, and C—in which you rationally invest equal credence. On hypothesis A, your evidence supports the belief that P, but on hypotheses B and C, it does not. A is true—your evidence really does support the belief that P. Given those stipulations, if you have all the beliefs supported by your evidence, then you'll find yourself believing that P (since A is true, believing that P is supported by your evidence) while believing that you probably shouldn't believe that P (since you think it more likely than not that hypothesis B or C is true, and according to those hypotheses your evidence does *not* support belief in P).

At this stage there are a wide variety of responses in the literature. Some authors simply bite the bullet, accepting that these sorts of epistemically akratic states are unavoidable—no adequate epistemological theory can avoid commitment to their possibility.¹⁶ Others hold that while there's *something* to be said for the akratic states, there's also something to be said against them—perhaps the sorts of cases just discussed present a kind of dilemma, in which not all epistemic demands can be met.¹⁷ And others try to contest the assumptions that led to the problem. Weaker versions of this third response will propose principles constraining just what combinations of first-order and higher-order evidence are possible, in the hopes that they can allow for *some* non-trivial HOE, but without licensing the sorts of seemingly irrational akratic states I've been discussing.¹⁸ A very strong version of this third response will involve claiming that agents

¹⁵The limited discrimination example in the previous section could be such a case, with A, B, and C corresponding to [51, 53], [52, 54], and [53, 55], and P being “the hand is not pointing to 54.”

¹⁶E.g., Aarnio (2014).

¹⁷E.g., Christensen (2010a) holds that subjects with misleading higher-order evidence can't help but violate a rational ideal. Worsnip (2015) holds that such subjects cannot both be rational, and believe what their evidence supports.

¹⁸Versions of this response include Elga (2013), Horowitz (2014), and Dorst (Ms.).

can neither be rationally uncertain of what their evidence is, nor of what it supports.¹⁹ Before presenting a version of the third response, I'll turn to a different set of puzzles.

3.2 Diachronic Puzzles

In addition to supporting incoherent belief states at a given time, non-trivial HOE leads to odd consequences concerning changes of evidence and belief over time. In this section I'll focus on odd diachronic consequences of ignorance concerning what your evidence *is*, rather than what a given body of evidence supports. But the latter phenomenon has been discussed in the literature too.²⁰

3.2.1 Monty Hall²¹

One of three doors, A, B, or C, is hiding a prize. The other two hide trash. Monty instructs you to pick a door. Before you can open it, Monty picks one of the remaining doors that he knows not to conceal a prize—there's guaranteed to be at least one—and opens it, revealing trash. He now gives you the choice of sticking with your initial choice, or switching to one of the remaining doors (either the open door with trash, or the closed door that might contain a prize).

For familiar reasons, you should switch to the closed door. There are many explanations, but the one I prefer runs as follows. Suppose you initially picked the wrong door—one that does not conceal the prize. Then there are two doors left, one of which contains the prize, and one of which does not. Given Monty's strategy, he'll have to

¹⁹Titelbaum (2015) is explicit that subjects cannot be rationally uncertain about evidential support relations, but doesn't discuss uncertainty about what one's evidence is. Salow (2017a,b) comes close—he argues against the possibility of rational uncertainty about evidential support relations, and argues against the possibility of many (but not all!) forms of rational uncertainty about what one's evidence is.

²⁰See, e.g., Horowitz (2014, §4.3).

²¹This section is inspired by Bronfman (2014), who notes the oddity of combining (1) the view that updating properly proceeds by conditionalization, (2) the view that KK failures are possible, and (3) the standard diagnosis of the Monty Hall problem. His solution is to hold onto 2 and 3, and propose a novel, non-conditionalization based form of updating. My aim in the present section is to make the sorts of KK failures that would be involved in Monty Hall cases seem odd enough that we'd be better off sticking with (1) and (3), and abandoning (2).

Fragmentation and Higher-Order Evidence

open the door that does not contain the prize, leaving the door that does contain the prize as the only one left. If you switch to the closed door, you'll switch to the door with the prize. Since you'll initially pick the wrong door two thirds of the time, switching will lead to the prize two thirds of the time.

Let's recast that explanation in terms of confidence. For each door, you start out with a credence of $\frac{1}{3}$ that the prize is behind that door. Suppose your initial pick is A. Your credences are unchanged—you have a credence of $\frac{1}{3}$ that the prize is behind door A, and $\frac{2}{3}$ that it's behind one of the other doors. Now, Monty reveals that the prize is not behind door C. Given what you know about Monty's strategy, you retain your credence of $\frac{2}{3}$ that it's behind door B or C, but now that you know it's not behind door C, all of that credence goes to B—you end up with a credence of $\frac{2}{3}$ that the prize is behind B, and $\frac{1}{3}$ that it's behind A.

So far, so familiar. But what if you could learn that the prize isn't behind door C, without learning that Monty revealed this to you? Would that undermine the argument for switching? Of course, there are *some* ways this could happen such that it obviously would. The door might be blown open by a gust of wind, not by Monty's design. Or you might have planted a hidden camera in the room, or hired a spy to look behind door C and report back. In cases like these, when you learn via a non-Monty source that the prize isn't behind door C, you have no reason to switch from A to B. But suppose you know for sure that the only way you could find out what's behind door C is by Monty's showing you, and you know how Monty operates (i.e., you know that he'll only open door C if the prize isn't there). In *this* case, is it still possible to learn that the prize isn't behind door C, without learning that Monty revealed this to you? And if so, would that undermine the argument for switching?

It's very hard to imagine how this could be. And yet, if we think that cases of limited discrimination provide illustrations of how one can gain evidence, without its becoming part of one's evidence *that* one has gained evidence, then it's hard to rule this out. E.g.,

Fragmentation and Higher-Order Evidence

suppose Monty operates as follows. Rather than throwing a door wide open to view, he renders a door ever so slightly transparent. Your perceptual discriminatory capacities are powerful enough, based on looking at the ever so slightly transparent door C, to give you the evidence that the prize isn't behind it—you can just barely discriminate the look of trash through a slightly transparent door from the look of a prize through a slightly transparent door. So when Monty renders the door slightly transparent, it becomes part of your evidence that the prize isn't behind door C. But your powers of discrimination are limited enough that it doesn't become part of your evidence *that* it's part of your evidence that the prize isn't behind door C.

If this is possible, how should you reason in such a case about whether to stick with door A or switch to B? You certainly aren't in a position to justify switching via the argument sketched earlier. That argument went via the premise that *Monty revealed* that the prize isn't behind door C, i.e., that he allowed you to learn this. But in the imagined case, that's not part of your evidence—you learned that the prize isn't behind door C, but you didn't learn *that* you learned this. The following reasoning sounds borderline incoherent, but it's hard to see what's wrong with it, given the assumption that limited discriminatory capacities can let you gain evidence without gaining evidence that you gained that evidence:

Clearly I shouldn't switch to C, because it doesn't contain the prize, but should I stick with A, or switch to B? If I *knew* C didn't contain the prize, then switching to B would be the sensible thing to do, because I'd only know that C didn't have the prize if Monty had revealed it to me, and a familiar argument would establish that I'd have a $\frac{2}{3}$ of gaining the prize by switching to B...but I can't tell whether I know that C doesn't contain the prize. For that matter, I can't tell whether I know that B doesn't contain the prize. Well, C doesn't contain the prize, so I guess I'll flip a coin to choose between A and B.

The next example, due to Bernhard Salow, has a similar flavor—it concerns odd consequences of allowing that one can gain evidence, without gaining evidence *that* one has gained evidence.

3.2.2 No-Lose Investigations

It's a commonsense piece of epistemology that an investigation can only provide support for a hypothesis if, had the results of the investigation been different, it could have undermined that hypothesis. Karl Popper (1963) famously criticized research programs that seemed to be able to accommodate any possible body of data; to use two of his favorite examples, if Freudian theories of psychodynamics or Marxist theories of history cannot be undermined by evidence, they cannot be supported by evidence either. While Popper's particular development of this insight was problematic,²² similar ideas about “no-lose” investigations have been defended in the recent epistemological literature.²³ But, as Salow (2017b) explains, once we take on the idea that you can gain evidence without gaining evidence that you gained evidence, we seem to be committed to the possibility of various sorts of no-lose investigation. Here's his explanation of how this works in Williamson's unmarked clock case, described earlier in §2.1.2:

My friend knows whether I'm popular; and I would like to have additional evidence that I am, regardless of whether it is true. So I construct an unmarked clock of the kind Williamson describes, and I ask my friend to set the hand in the following way: if people like me, he will set it to 53; if they don't, he will flip a coin to decide whether to set it to 52 or to 54. Having given the instructions, I know that the clock will be set somewhere between 52 and 54...

Next, I take a look. If people actually like me, the hand will be set to 53,

²²See Salmon (1981).

²³See, e.g., Titelbaum (2010), Salow (2017b).

Fragmentation and Higher-Order Evidence

and so my evidence will only tell me that it is somewhere between 52 and 54, which I knew already. So if people like me, I get no new evidence. But if people do not like me, it will be set either to 52 or to 54. Suppose it is set to 52; then my evidence will allow me to rule out that it's set to 54, since 54 is far enough away from the actual setting. But I knew that there was a fifty-fifty chance that it would be set to 54 if people didn't like me. So seeing that it isn't set to 54 gives me some evidence that I am popular. Moreover, my evidence cannot discriminate between the hand being set to 52 and its being set to 53, so that I get no evidence against my being popular. So, if the hand is set to 52, I will get evidence that I am popular; by similar reasoning, I will also get such evidence if the hand is set to 54. So if people don't like me, I will get evidence that I am popular. Again, I have successfully set up a no-lose investigation into my popularity. (pp. 18-19)

If we ask how the subject of Salow's example should reason, or might rationally act on his evidence—imagine that he's deciding whether to book a large venue for his birthday party, which would only be necessary if he's popular—we'll run into the same sorts of perplexities we saw in the case of Monty Hall.

I don't pretend to have provided anything like a comprehensive overview of the terrain. I've just tried to point to various sorts of awkwardness that ensue when we allow that subjects can be rationally ignorant of either (a) which body of evidence is theirs, or (b) which beliefs are supported by which bodies of evidence. But the idea that subjects *can't* be rationally ignorant of these things can sound hard to swallow. So in the next section I'll present a framework that let's us capture the truths that we express when we describe subjects as being rationally ignorant of (a) and (b), but without the awkward consequences.²⁴

²⁴Does it have *other* awkward consequences? Certainly.

4 Fragmented Evidence

What does it take to have some proposition as part of your evidence? First, assume the debatable but defensible view that $E = K$ —your evidence is your knowledge.²⁵ Also assume that belief is a necessary condition on knowledge—you only know that P if you believe it. In this section I’ll try to show how broadly “fragmentationist” views about belief can lead to similarly fragmentationist views about evidence, and I’ll then apply those views to the putative cases of HOE discussed earlier in this essay.

What does it take to believe that P ? A debatable, but attractive answer is that it involves behaving in ways that make sense if P . E.g., someone who believes that there’s beer in the fridge will be disposed, all else equal, to walk to the fridge if she wants a beer, to respond “yes” to the question “is there any beer left?”, to offer beer to thirsty guests, to be surprised upon opening the fridge and failing to see beer, etc.—these are behaviors that make sense if there’s beer in the fridge.²⁶ But one can act as if P in some situations but not others, or when engaged in certain tasks but not others. E.g., one might *talk* like a P -believer, but *walk* like a P -disbeliever.²⁷ We might capture this by saying that, for the purpose of talking, you believe that P , but for the purpose of walking you believe that $\sim P$ —your doxastic life is fragmented. And we can cut more finely. Even within the category of verbal behavior, or non-verbal behavior, some of one’s actions might make sense if P , but not others.²⁸

This kind of fragmentationist view about belief leads naturally to a fragmentationist

²⁵Famously defended by Williamson (2000).

²⁶The example is from Schwitzgebel (2002), who defends a view about belief broadly similar, I think, to the one sketched in this paragraph. There he says believing that P involves “fitting the dispositional stereotype” of someone who believes that P . While I’m inclined to think that’s right, I’m more optimistic than he is about the prospects for giving a somewhat systematic account of which dispositions get stereotypically associated with which beliefs, along broadly radical interpretationist lines, where we can say something systematic about what sorts of dispositions “make sense” for a subject who believes that P . See Lewis (1974), Stalnaker (1984), and Dennett (1987). Also, unlike the previous authors, Schwitzgebel stresses that he thinks dispositions to have irreducibly phenomenal states are among the dispositions associated with having various beliefs. For my purposes in this essay, I don’t think I need to take a stand on that question.

²⁷See, e.g., Schwitzgebel (2001).

²⁸See, e.g., Elga and Rayo (Manuscript A, B).

Fragmentation and Higher-Order Evidence

view about knowledge and (if $E = K$) evidence. That is, suppose S has a dispositional profile that amounts to believing that P for some purposes, but not others. And assume that P is true, and that S 's (fragmentedly) believing that P has an appropriately non-accidental connection to the truth of P —the sort of connection necessary for knowledge. Then it's natural to say that S knows that P —and thereby has P as part of her evidence—for some purposes, but not others. Jack Marley-Payne (Ms.) offers a nice example that illustrates the idea:

Take the example of the inarticulate tennis player—let's call her Serena. Over the course of a rally she can execute a complex plan which involves hitting repetitive shots to first ground her opponent in one position and then wrong foot him in order to win the point. Moreover, she can calibrate her play in response to the court conditions, the abilities of her opponent, whether she desires to humiliate him etc. However, she is unable to explain what she was doing—indeed she may even say things about her play that turn out to be false. Her non-verbal behaviour suggests a belief that a cross-court shot will wrong foot her opponent while her verbal behaviour does not.

Suppose that a cross-court shot *will* wrong foot her opponent, and it's no accident that her non-verbal behavioral dispositions reflect sensitivity to this fact; rather, it's the product of years of training, which included seeing how opponents reacted to similar situations in the past. In that case, it's natural to say that Serena *knows* that a cross-court shot will wrongfoot her opponent. While she's unable to articulate her knowledge, this doesn't relegate it to the realm of mere reflex; after all, her choice of a cross-court shot is sensitively dependant on her background beliefs and desires in a way that mere reflexes aren't. If we think the characteristic functional role of beliefs is to combine with other beliefs to lead you to choose actions that will satisfy your desires—while mere reflexes manifest themselves in rigid ways that aren't so sensitive to the rest of an agent's psychology—then what Serena has is belief, rather than mere reflex. And because

that belief is non-accidentally true, it's knowledge. It just happens to be knowledge that she can't draw on for verbal reports.²⁹

We finally have enough conceptual machinery on the table to start applying fragmentation to the topic of higher-order evidence. With the fragmentationist view of belief and evidence in hand, we can distinguish two readings of the claim that non-trivial HOE is impossible. Let x and y be purposes relative to which a subject might have evidence.

Fixed Purpose: It's never the case that S 's evidence $_x$ supports the claim that P , while her evidence $_x$ supports the claim that her evidence $_x$ doesn't support the claim that P .

Variable Purpose: It's never the case that S 's evidence $_x$ supports the claim that P , while her evidence $_y$ supports the claim that her evidence $_x$ doesn't support the claim that P .

In the remainder of the paper, I'll argue that **Fixed Purpose** is what we need to rule out the rationality of epistemic akrasia, and the bizarre verdicts about Monty Hall and the possibility of no-lose investigations. But we don't need **Variable Purpose**, and familiar putative examples of non-trivial HOE can be interpreted as counterexamples to it. My hope is that this amounts to having our cake and eating it too. We can agree with a version of the intuitively plausible idea that you're not always in a position to tell

²⁹The idea that what you know depends in some sense on what task you're engaged in is reminiscent of the version of "contextualism"—in contemporary terminology, probably best interpreted as a version of sensitive invariantism—defended by Michael Williams (1977, 1991). He defends a view on which our inability to answer skeptical challenges amounts to lacking knowledge, but only while we are engaged in the epistemological project of answering skeptical challenges. When we turn our attention to other inquiries—history, ornithology, or just what to have for lunch—our knowledge returns. While I'm sympathetic to much of Williams' discussion, a crucial difference between his approach and the present one concerns higher-order knowledge and evidence. He understands his view as crucially connected to the possibility of failures of KK—cases in which a subject knows, but fails to know that she knows. E.g., "Thus stated, contextualism implies a kind of externalism, for though appropriate contextual constraints will have to be met, if a given claim is to express knowledge, they will not always have to be known, or even believed, to be met." Williams (1991, p.119) "Externalism thus drives a wedge between knowing something or other and knowing that one knows it." (p. 96). By contrast, in the remainder of this paper I'll argue that fragmentationist views about knowledge and evidence open up room for reinterpreting cases that seemed to involve straightforward divergences between first-order and higher-order epistemic status.

Fragmentation and Higher-Order Evidence

what your evidence supports, and we can apply that idea in broadly the ways it's been applied in the literature. But by distinguishing one version of that idea from a closely related one, we can explain why, properly interpreted, it doesn't lead to absurdity.

I'll start with the case of the unmarked clock from §2.1.2. On the interpretation I'll suggest, the subject of the example has different information available for different purposes. On my way of fleshing out the case, the subject is guided by one body of information in her assertions and judgments, but a different body of information in her visuomotor behavior. If we add the subscripts $_{aj}$ and $_{vm}$ to "evidence", so that the subject's $evidence_{aj}$ is the knowledge available for guiding assertions and judgments, and her $evidence_{vm}$ is the knowledge available for guiding visuomotor tasks, then the diagnosis will be as follows: the subject's $evidence_{aj}$ fails to settle what her $evidence_{vm}$ is. But nothing in the case suggests that her $evidence_{aj}$ fails to settle what her $evidence_{aj}$ is, and, as I'll explain, that's enough to avoid awkward consequences concerning Monty Hall, or no-lose investigations.

Next, I'll turn to cases of putative ignorance about evidential support relations, such as the JUROR case from §2.1.1. I'll argue that a fragmentationist can make sense of the idea that one ought to be agnostic about whether the accused is at fault, but without understanding that verdict as straightforwardly conflicting with the requirement of total evidence, and without the more general threat of rational akrasia posed by limited evidence concerning epistemic support relations. The basic idea will be that, in an important sense, the relevant evidence the subject has available for determining liability is a relatively sparse one, rather than the richer body that includes all of the evidence that was presented at court. The subject has that richer body of evidence in a *sense*, but only a sense—she has pieces of that body of evidence available for various different purposes, but she doesn't have the *whole* body of evidence available for the purpose of determining liability.

4.1 Limited Discrimination Revisited

In Salow’s description of the unmarked clock case, which mirrors Williamson’s, we’re told that when the clock’s hand is in fact indicating that it is 53 minutes past the hour, your powers of discrimination enable you to rule out that it indicates 51 minutes or fewer, and also to rule out that it indicates 55 or more, but they don’t allow you to discriminate any more finely than that—you cannot rule out 52, 53, or 54. What form does this discrimination take—how do your powers of perceptual discrimination manifest themselves? While this question isn’t given much attention by either Williamson or Salow, I’ll try to show that it’s in fact of central importance.

A natural answer is that they manifest themselves in occurrent judgments, and/or assertions—when the clock is in fact pointing at 53, then you will be disposed to inwardly judge, and/or outwardly assert, claims like the following: “it’s not at 51 or lower,” “it’s not at 55 or higher”, “it might be at 52, 53, or 54.”³⁰ And likewise, more generally, for other positions—when the clock is in fact pointing towards n , you’ll be disposed to judge and/or assert that it’s not at $n - 2$ or below, not at $n + 2$ or above, but nothing stronger than that.

So far, so good. However, this way of thinking about how her powers of discrimination manifest themselves is inconsistent with another crucial feature of the case, as Williamson and others who’ve written about it understand it.³¹ Namely, you yourself—the subject of the case—could know the general epistemic features of the case; you can know that when the clock is pointing to n , your powers of discrimination enable you to know that it is pointing between $n - 1$ and $n + 1$, but nothing stronger. Intuitively, this is plausible—you could have reliable information about just how good your vision is, perhaps provided by an optometrist. Why is the idea that your powers of discrimination manifest themselves

³⁰While Williamson mainly describes the case in terms of “discrimination” and “knowledge” without indicating how they might be manifested, Salow explicitly talks about what the subject can “reliably judge.” (p. 17)

³¹See especially (Christensen, 2010b, p.133-4).

in judgment and assertion inconsistent with the idea that you can know the general epistemic features of the case? For the following reason. Suppose the clock is pointing at 53. Then your powers of perceptual discrimination will manifest themselves in your making the assertions and judgments mentioned in the previous paragraph—explicitly judging and/or asserting that it might be between 52 and 54, but nowhere else. But if you make those judgments or assertions, you can also notice that you’ve made them. And making those assertions and judgments is only compatible with knowing that it is between 52 and 54—on the assumption that your perceptual knowledge manifests itself in judgment and assertion, then if you had different perceptual knowledge you’d make different assertions and judgments. But then, contra the stipulation of the case, you’re in a position to know what your evidence is—namely, that the clock is pointing between 52 and 54. And because you know the general epistemic features of the case, we get the absurd result that you are in a position to know that the clock is pointing to 53—since you know that your evidence is only [52, 54] when it’s *actually* pointing to 53.

What went wrong? One quick response is just to deny that it’s possible for there to be any version of the case in which both (a) your powers of perceptual discrimination are limited, and (b) your visual evidence is perfectly sensitive to the actual position of the clock. To motivate this response, we might say that the thought that your powers of discrimination are limited—i.e., that you can’t tell exactly where the clock is pointing just by looking—is intimately bound up with the idea that, compatible with your having the visual evidence you have, there are various different positions the clock might be pointing. So if we want to hold on to the idea that you have limited powers of discrimination, we should reject the idea that there’s a 1-1 function from positions of the clock to bodies of evidence you get when you look at the clock. This is the route taken by Stalnaker (2009) in the course of responding to a similar case in Hawthorne and Magidor (2009). Applying his response to the present case would lead to the result that, rather than being determined by the actual position of the clock, your evidence is determined

Fragmentation and Higher-Order Evidence

by your *best guess* as to the position of the clock, which you are in a position to know. Moreover, there's no strict mapping from where the clock is pointing to what your best guess will be. So the question "what is the subject's evidence when the clock is in fact pointing to position n ?", will have a range of possible answers, rather than a unique one.

While I'm sympathetic to this response, I think we can be a bit more concessive to the intended interpretation of the case. That is, we can allow that there's a sense in which the subject's evidence depends on the actual position of the clock, which the subject is nevertheless not in a position to know, while still resisting the central lesson about higher-order ignorance. We were led into absurdity when we assumed that your powers of perceptual discrimination would manifest themselves in a way that you were in a position to notice, as when you notice what you say or judge. So suppose they don't. Suppose that your powers of perceptual discrimination manifest themselves in your visuomotor behavior—e.g., they manifest themselves in the direction your hand will move if you decide to reach for the pointer, and other tasks that require integrating visual information with motor activity. For example, suppose your hand-eye coordination isn't perfect, and if asked to very quickly reach out and touch the clock's pointer, you won't always hit it (imagine the hand is quite thin, and you're asked to move very quickly). However, when the pointer is in fact pointing at n , you'll never touch a position below $n - 2$, or above $+2$. And suppose this isn't just a quirky fact about a particular task—for a wide range of visuomotor tasks, when the clock is pointing at n , your visuomotor behavior will be as if the clock is pointing somewhere in $[n - 1, n + 1]$.

This lets us vindicate *part* of the description of the case—the part concerning how your evidence about where the clock is depends on where the clock is actually pointing. But what about the description of the case in terms of *higher-order* evidence? That is, can we make sense of the idea that, when the clock is pointing at 53, your evidence leaves open various possibilities concerning what your evidence is? If we stick to the idea that what evidence you have is manifested in your visuomotor behavior, it's hard

Fragmentation and Higher-Order Evidence

to see how to do so. That is, while it's not so hard to see how, insofar as we're interested in explaining your visuomotor behavior, we might fruitfully interpret you as having the evidence that the clock is pointing somewhere in [52, 54], it's much harder to see how, insofar as we're interested in explaining your visuomotor behavior, we might fruitfully interpret you as lacking evidence about what your evidence is. That is, it's just not clear how we should expect information or lack thereof *concerning what your evidence is* to be manifested in visuomotor behavior.

So let's return to assertion and judgment. Plausibly, when you are in fact completing visuomotor tasks as if the clock is pointing in [52, 54], you won't be in a position to reliably assert or judge that you are doing so. E.g., suppose the clock is pointing to 53, and you've just been asked to reach out and touch the pointer 10 times in a row. Each time you touch a position somewhere in [52, 54]. Now you're asked: "did you touch a position below 52 any of those times?" Plausibly, you'll respond: "I don't know, maybe. My best guess is that the clock is pointing to 53, but that's just a guess, so it might be pointing to 52. And given that my hand-eye coordination isn't perfect, if it's in fact pointing to 52, I probably hit 51 a few times." And your coherently giving that speech is consistent with your being happy to assert that *if* the clock is in fact pointing at 53, then you did not touch a position below 52 on any of those occasions. Unlike when we were thinking of the manifestation of one's knowledge solely in terms of assertion or judgment, now we can see how one's knowledge that the clock is pointing in [52, 54] could be reliably manifested, without your being in a position to learn from that reliable manifestation *that* you know that the clock is in [52, 54].

To keep all this straight, it will be helpful to introduce subscripts. Let your evidence_{e_{aj}} be your knowledge available for guiding assertions and judgments, and your evidence_{e_{vm}} be the knowledge available for guiding visuomotor tasks. The diagnosis I've been offering can be stated as follows. When the clock is in fact pointing at n , your evidence_{e_{vm}} is that it is pointing somewhere in $[n - 1, n + 1]$. Moreover, that general fact about

how your evidence_{vm} sensitively depends on the actual position of the clock is part of your evidence_{aj}. And when your evidence_{vm} is that the clock is pointing somewhere in $[k - 1, k + 1]$, your evidence_{aj} doesn't settle that this is your evidence_{vm}—rather, there will be various different possibilities compatible with your evidence_{aj} concerning what your evidence_{vm} is. Which will those be? Here, we cannot give a unique answer, because while your evidence_{vm} sensitively depends on the actual position of the clock, your evidence_{aj} does *not*—that way lies absurdity, as we saw. What *does* your evidence_{aj} depend on? Here I'm happy to stick with Stalnaker's model—perhaps it depends on your “best guess” as to the position of the clock.³²

The present diagnosis lets us avoid awkward consequences concerning Monty Hall. We can get a version of the Monty Hall case going where one “learns” that the prize isn't behind door C—in the sense that one's visuomotor dispositions now involve being disposed to complete visuomotor tasks as if it's not there, in some way non-accidentally connected to the fact that it's really not there—but where one doesn't gain this information in a form that would make it available for assertion, judgment, and crucially, deliberation. And since the apparent absurdity in the case involved trying to imagine the sort of planning and deliberation that would be appropriate for an agent who had the evidence that the prize wasn't behind door C, but lacked the evidence *that* she had this evidence, that absurdity vanishes when the only sense in which she “has” the evidence that the prize isn't behind door C is a sense that wouldn't be manifested in her deliberation. The diagnosis of Salow's “no-lose inquiry” case would proceed along much the same lines.

Taking a step back, in what sense does my diagnosis do without non-trivial HOE? It's

³²As I hope is clear enough, it's not central to my diagnosis that it's *visuomotor* information that we contrast with information available for guiding assertion and judgment. It's just that there be *some* such contrast—some type of task that one's evidence (in a sense) can be used to guide, but whose manifestation can't be reliably noticed and remarked upon by the subject who has that evidence. And since assertions and judgments *can* be noticed and learned from—when you assert or judge that P, you're in a position to know that you've asserted or judged that P—for a diagnosis of my sort to work, we need some way that knowledge can be manifested other than in assertion and judgment. Visuomotor behavior is just one convenient, natural alternative.

true that there's a sense in which my explanation of the case makes appeal to non-trivial evidence concerning what your evidence is. But the two occurrences of "evidence" in the previous sentence require different readings—that is, I say that in the case of the unmarked clock, you have non-trivial evidence_{aj} concerning what your evidence_{vm} is. But if we hold fixed the subscript, we don't have non-trivial higher-order evidence: we saw that non-trivial evidence_{aj} concerning one's evidence_{aj} led to absurdity, and it's not clear how to interpret talk of evidence_{vm} concerning one's evidence_{vm}. So while my diagnosis lets us say that in the case of the unmarked clock, your evidence doesn't settle what your evidence is, that turns out to be a less than maximally perspicuous description of the situation. And my suspicion is that this will hold quite generally—cases that are tempting to describe as involving a subject's evidence failing to settle what her evidence is will always turn out to be cases where we can distinguish different senses of "evidence", and so distinguished, they will be cases in which one's evidence—in *one* sense—fails to settle what one's evidence—in some *other* sense—is. Given the distinction drawn earlier, they will be counterexamples to **Variable Purpose**, but not **Fixed Purpose**. While I can't offer a full defense of that claim here, I hope I've at least shown the following: absent some such distinction, we get puzzles and paradoxes, as illustrated in §3.2. And in a general, schematic sort of case meant to illustrate the possibility of rational ignorance concerning one's evidence—the case of the unmarked clock—the diagnosis I've offered is a natural and attractive one that avoids the paradoxes. So it's natural to hope that similar diagnoses should be available in other cases.

4.2 Epistemic Support Relations Revisited

What about the cases meant to illustrate higher-order uncertainty not via ignorance about what one's evidence *is*, but instead about what it *supports*? Can the strategy of the previous subsection be of any help? I believe it can, though I admit the fit here is a bit less natural.

There is a constellation of closely related suggestions that a number of writers have defended for how to deal with HOE that prompts uncertainty about epistemic support relations—sometimes it’s put in the language of “bracketing” (Elga, 2007; Christensen, 2010a), sometimes “calibration” Sliwa and Horowitz (2015); Schoenfield (2015), sometimes in other terms.³³ The suggestion that we should use these methods is typically treated as a kind of *sui generis* epistemological principle, which may conflict with other more general principles—most notably, the requirement of total evidence (Kelly, 2010). My ultimate goal in this section is as follows. I want to show how the kind of “bracketing” recommended by the authors just mentioned fits naturally with the idea, defended in the previous section, that evidence can be available for guiding some tasks, but not others. Since this idea is plausible and defensible independently of considerations having to do with higher-order evidence, my aim is to make “bracketing” seem more principled, and less ad hoc. I also hope that my discussion will render the *prima facie* conflict with the total evidence principle more palatable, in part by making it a bit less clear just how the total evidence principle should be interpreted, once we adopt the fragmentationist view of evidence.

Return to the case of JUROR, described in §.2.1.1. The recommendation of the bracketer is that, after hearing about the disagreement of her peers, the juror should form her opinion as follows. She should set aside the particular details of the evidence presented at trial and why she thinks those details point to the liability of the accused. Instead, she should reason from the sparse body of essentially psychological evidence, which includes only (a) the fact that she formed the judgment she did, (b) the fact that her peers formed the judgments they did, and (c) general facts that bear on her and her peers’ reliability in matters like these. And that sparse body of evidence will support, I assume, agnosticism about whether the accused is at fault.

This recommendation looks like it straightforwardly conflicts with the requirement

³³Roger White (2009) calls it “treating oneself and others as thermometers.” In this volume, Sophie Horowitz discusses a similar suggestion couched in terms of “perspective.”

of total evidence.³⁴ There are facts that are part of the juror’s total evidence, and which bear on the question of whether the accused is at fault, that the bracketer says should nevertheless be set aside or ignored when reasoning about whether the accused is in fact at fault.

In light of the previous section, however, I believe we should be a bit uneasy when we see phrases like “the juror’s total evidence”. After all, we saw there that it can be helpful to distinguish between the evidence a subject has available for some tasks from the evidence she has available for others. And in such cases, “the subject’s total evidence” will be threaten to be ambiguous, or liable to refer to different bodies of evidence depending on which task we are contemplating the subject performing. While in that section we only distinguished between broad sorts of tasks—discursive versus visuomotor—there’s no principled reason we can’t distinguish more finely. Elga and Rayo (Manuscript A, B), following up on some suggestions of Stalnaker (1991, 1999), argue that solving the problem of logical omniscience requires distinguishing the information a subject has for verbally answering some questions from information a subject has for verbally answering others.³⁵

How does this possibility bear on the case of JUROR? Here’s the idea. While the individual items of evidence presented at trial are each available to the juror for various purposes—answering pointed questions about whether, e.g., the accused drove a sedan, or whether the plaintiff was in Tallahassee on March 21—they may not be collectively available for answering the question: “is the accused at fault?” For that latter task, perhaps the only evidence available is the sparse body of broadly psychological evidence that, according to the bracketer, should determine the juror’s view.

Phenomenologically at least, this strikes me as plausible. When dealing with a simple body of evidence, I feel like I can base my beliefs directly on that evidence—e.g., if I see

³⁴See Williamson (2007) on the dangers of “psychologizing the evidence.”

³⁵While not put in quite these terms, Rayo (2011) is in a very similar spirit, and nicely complements my strategy in the text for addressing the JUROR case.

Fragmentation and Higher-Order Evidence

that the streets are wet, I'll straightaway think that it recently rained, and it's clear to me how my evidence supports my belief. But when dealing with a sufficiently complex and multifaceted body of evidence, with different pieces pointing in different directions and in ways that can't be straightforwardly weighed against one another, the process feels quite different. I'll think about it for awhile, and eventually I'll find myself stably inclined to think one thing or another, but not in a way where it's at all transparent to me how the particular pieces of evidence combine to generate my resultant doxastic inclination. Suppose I accede to that inclination—e.g., I'm in the position of JUROR before having heard from my peers, I find myself inclined to believe that the accused is probably guilty, and I go on to believe that. I'm tempted to say that in a case like this my ultimate belief is based not directly on my evidence—as it is in the simple cases—but is instead based on an inclination prompted by reflection on that evidence. While the individual items of evidence are available to me *in a sense*—I can base *some* beliefs directly on them, taken individually or in small groups—they're not collectively available to me for answering the target question.

I didn't mention disagreement in the previous paragraph. But it's not hard to factor in. Just as I can base a belief that the accused was at fault on the fact that I was inclined to so believe that after reflecting on the evidence, I can base agnosticism about whether the accused was at fault on the facts that (a) I was initially inclined to believe that the accused was at fault, (b) my peers were inclined to believe the opposite, and (c) none of have any distinctive advantage when it comes to evaluating evidence of this sort.

I admit that this view complicates the total evidence requirement, and threatens to reduce it to triviality—if a subject can always avoid the charge of violating the requirement by saying that the total evidence wasn't really “available” for the task in question, then what does the requirement rule out? But rather than a reason to reject the strategy I've been outlining for how to think about disagreement, I think this is a genuine difficulty with how to interpret the total evidence requirement once we allow for the

Fragmentation and Higher-Order Evidence

possibility of fragmented belief and evidence, which I think we have independent reason to do. We don't want the total evidence requirement to rule that a subject who lacks the ability to integrate the information available to her for guiding visuomotor tasks with the information available for guiding explicit verbal reasoning to thereby count as irrational; there's nothing irrational about being able to hit a bullseye without being able to explain what you're doing. But once we admit that, it becomes difficult to say just what sorts of failures to integrate bodies of information available for distinct tasks count as failures of rationality, of the sort that the total evidence requirement rules irrational. And absent an answer to that question, it's not clear whether, properly interpreted, the total evidence requirement will rule out the kind of "bracketing" that's been discussed in the literature on disagreement and HOE more generally.

How broadly will this strategy work, if it works at all, for vindicating "bracketing" as a response to HOE prompting uncertainty about evidential support relations? My diagnosis of JUROR depended on the idea that the subject in question is dealing with a complex body of evidence whose significance for the target question is difficult to discern, but not all putative cases of rational uncertainty about epistemic support relations have that structure. E.g., what about a case where my evidence is that the streets are wet, and I'm inclined to believe that it rained on that basis, but I then find out that I've taken a drug that makes me very bad at evaluating evidence?³⁶ This case, I think, is much harder to fit into the mold of the previous one, since it's not clear in what sense I might fail to be in a position to base a belief that the streets are wet directly on the fact that it rained, even after taking the drug. I mention the case only to set it aside—I myself have a very hard time wrapping my head around such cases, and am inclined to think not much weight should be put on them. It's not clear to me what we should say about what evidence is available to such a subject for what purposes, nor is it clear to me what we should say such a subject should believe in such a case. While I feel

³⁶This case is inspired by Christensen (2007a).

Fragmentation and Higher-Order Evidence

the force of the bracketer's position in cases like JUROR, I think the fragmentationist can nicely handle those cases. I'm not sure what the fragmentationist can or should say about cases like this one, but that uncertainty matches my uncertainty about what the *right* verdict is.

What about the description of disagreement cases in terms of HOE? It seemed natural to say that the juror's evidence supported a particular belief, but failed to support the claim *that* it supported that belief. How should we revise this claim if we're fragmentationists about evidence, along the lines I've been discussing? The closest thing we can say, I think, is the following. The evidence the juror *would have* for the purpose of determining whether the accused was at fault, *if* she had cognitive powers far vaster than ours—vast enough for the inference from that evidence to the accused's fault to be as transparent and direct for her as the inference from wet streets to rain is for us—supports the claim that the accused is at fault. But the evidence the juror *in fact* has available for the purpose of determining whether the accused was at fault fails to support the claim in the previous sentence—it fails to support the claim that the evidence she *would have* if she had with more impressive cognitive powers supports the claim that the accused is at fault. So there's no single reading of "evidence" on which the juror's evidence supports P, while also supporting agnosticism concerning whether it supports P—no counterexample to **Fixed Purpose**. Rather, the sense in which it supports P, and the sense in which it supports agnosticism, are quite different.

And that lets us avoid the threat of epistemic akrasia. Once we offer the present diagnosis, we're under no pressure to say that the juror's evidence supports believing some claim, while also believing that she shouldn't believe it. Rather, we can say the evidence she has available for determining whether the accused is at fault unequivocally supports agnosticism on that question, and supports believing that she should be agnostic on that question. Of course, if the subject had more impressive cognitive powers, then she'd have a different body of evidence available for determining whether the accused is

at fault—one that would support an affirmative answer. But that’s different from her *already* having a body of evidence that would rationalize an affirmative answer, (while failing to rationalize the claim *that* it would rationalize an affirmative answer) and so we don’t need to say that her presently available evidence—for any purpose—supports an akratic state.

5 Conclusions

When we say without qualification that a subject’s evidence can fail to settle what her evidence supports—either because it fails to settle what her evidence *is*, or because it fails to settle questions about which bodies of evidence support which propositions—we run into troubles. But to flatly deny that this is possible beggars belief.³⁷ My strategy in this paper has involved following the old dictum: “whenever you meet a contradiction, draw a distinction.”³⁸ By allowing counterexamples to **Variable Purpose**—allowing *one* version of the idea that there can be non-trivial HOE—we can have a framework for describing a host of cases of epistemological interest. But by ruling out counterexamples to **Fixed Purpose**—rejecting a *different* version of the idea that there can be non-trivial HOE—we can avoid the troubles that the more familiar descriptions engender.

Bibliography

Aarnio, Maria Lasonen. 2014. “HigherOrder Evidence and the Limits of Defeat.” *Philosophy and Phenomenological Research* 88:314–345.

Bronfman, Aaron. 2014. “Conditionalization and Not Knowing That One Knows.” *Erkenntnis* 79:871–892.

Christensen, David, 2007a. “Does Murphy’s Law Apply in Epistemology? Self-Doubt

³⁷See Dorst (Forthcoming).

³⁸William James attributes it to the Scholastics.

Fragmentation and Higher-Order Evidence

- and Rational Ideals.” In Gendler, Tamar Szabó, and John Hawthorne, editors, “Oxford Studies in Epistemology, Volume I,” pages 3–31. Oxford University Press.
- . 2007b. “Epistemology of Disagreement: The Good News.” *Philosophical Review* 116:187–217.
- . 2010a. “Higher-Order Evidence.” *Philosophy and Phenomenological Research* 81:185–215.
- . 2010b. “Rational Reflection.” *Philosophical Perspectives* 24:121–140.
- Dennett, Daniel, 1987. *The Intentional Stance*. Cambridge, MA: The MIT Press.
- Dorst, Kevin, Forthcoming. “Higher-Order Uncertainty.” In “Higher-Order Evidence,” Oxford University Press.
- . Unpublished Manuscript. “Evidence: A Guide for the Uncertain.” .
- Egan, Andy. 2008. “Seeing and Believing: Perception, Belief Formation and the Divided Mind.” *Philosophical Studies* 140:47 – 63.
- Elga, Adam. 2007. “Reflection and Disagreement.” *Nous* 41:478–502.
- . 2013. “The Puzzle of the Unmarked Clock and the New Rational Reflection Principle.” *Philosophical Studies* 164:127–139.
- . Unpublished Manuscript. “Lucky to be Rational.” .
- Elga, Adam, and Agustín Rayo. Unpublished Manuscript. “Fragmentation and Information Access.” .
- Elga, Adam, and Agustín Rayo. Unpublished Manuscript B. “Fragmentation and Logical Omniscience.” .

Fragmentation and Higher-Order Evidence

- Feldman, Richard, 2007. "Reasonable Religious Disagreements." In Antony, Louise, editor, "Philosophers Without gods," pages 194–214. Oxford University Press.
- . 2009. "Evidentialism, Higher-Order Evidence, and Disagreement." *Episteme* .
- Gendler, Tamar Szabó. 2008. "Alief and Belief." *Journal of Philosophy* 105:634–663.
- Greco, Daniel. 2014a. "Iteration and Fragmentation." *Philosophy and Phenomenological Research* 88.
- . 2014b. "A Puzzle About Epistemic Akrasia." *Philosophical Studies* 167:201–219.
- Hawthorne, John, and Ofra Magidor. 2009. "Assertion, Context, and Epistemic Accessibility." *Mind* 118:377–397.
- Horowitz, Sophie. 2014. "Epistemic Akrasia." *Noûs* 48:718–744.
- . Forthcoming. "Predictably Misleading Evidence." In "Higher-Order Evidence," Oxford University Press.
- Kelly, Thomas, 2010. "Peer Disagreement and Higher Order Evidence." In "Disagreement," Oxford University Press.
- Lewis, David. 1974. "Radical Interpretation." *Synthese* 27:331–344.
- . 1982. "Logic for Equivocators." *Noûs* 16:431–441.
- Littlejohn, Clayton. 2018. "Stop Making Sense? A Puzzle about Evidence and Epistemic Rationality." *Philosophy and Phenomenological Research* .
- Marley-Payne, Jack. Unpublished Manuscript. "Task-Indexed Belief." .
- Popper, Karl, 1963. *Conjectures and Refutations*. Routledge.
- Rayo, Agustín. 2011. "A Puzzle About Ineffable Propositions." *Australasian Journal of Philosophy* 89:289–295.

Fragmentation and Higher-Order Evidence

- Rayo, Agustín, 2013. *The Construction of Logical Space*. Oxford University Press.
- Rinard, Susanna. Forthcoming. “Reasoning One’s Way Out of Skepticism.” *Brill Studies in Skepticism* .
- Salmon, Wesley. 1981. “Rational Prediction.” *British Journal for the Philosophy of Science* .
- Salow, Bernhard. 2017a. “Elusive Externalism.” *Mind* .
- . 2017b. “The Externalist’s Guide to Fishing for Compliments.” *Mind* .
- Schoenfield, Miriam. 2015. “A Dilemma for Calibrationism.” *Philosophy and Phenomenological Research* 91:425–455.
- Schwitzgebel, Eric. 2001. “In-Between Believing.” *Philosophical Quarterly* 51:76–82.
- . 2002. “A Phenomenal, Dispositional Account of Belief.” *Nous* 36:249–275.
- Sliwa, Paulina, and Sophie Horowitz. 2015. “Respecting All the Evidence.” *Philosophical Studies* 172:2835–2858.
- Smithies, Declan. 2015. “Ideal Rationality and Logical Omniscience.” *Synthese* 192:2769–2793.
- Stalnaker, Robert, 1984. *Inquiry*. Cambridge, MA: The MIT Press.
- . 1991. “The Problem of Logical Omniscience, I.” *Synthese* 89.
- , 1999. “The Problem of Logical Omniscience II.” In “Context and Content,” Oxford University Press.
- . 2009. “On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility.” *Mind* 118:399–409.

Fragmentation and Higher-Order Evidence

- Titelbaum, Michael, 2015. "Rationality's Fixed Point (Or: In Defense of Right Reason)." In Gendler, Tamar Szabo, and John Hawthorne, editors, "Oxford Studies in Epistemology, Volume 5," Oxford University Press.
- Titelbaum, Michael G. 2010. "Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology." *Philosophical Studies* 149:119–134.
- White, Roger. 2009. "On Treating Oneself and Others as Thermometers." *Episteme* 6:233–250.
- Williams, Michael, 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. B. Blackwell.
- , 1999—First edition published in 1977. *Groundless Belief*. Princeton University Press.
- Williamson, Timothy, 2000. *Knowledge and its Limits*. Oxford University Press.
- , 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.
- , 2011. "Improbable Knowing." In "Evidentialism and its Discontents," Oxford University Press.
- Worsnip, Alex. 2015. "The Conflict of Evidence and Coherence." *Philosophy and Phenomenological Research* 91.